arXiv:1608.01338v2 [cs.AI] 5 Aug 2016

Paraconsistency and Word Puzzles

TIANTIAN GAO, PAUL FODOR and MICHAEL KIFER
Stony Brook University, New York, USA
(*e-mail:* {tiagao, pfodor, kifer}@cs.stonybrook.edu)

Abstract

Word puzzles and the problem of their representations in logic languages have received considerable attention in the last decade (Ponnuru et al. 2004; Shapiro 2011; Baral and Dzifcak 2012; Schwitter 2013). Of special interest is the problem of generating such representations directly from natural language (NL) or controlled natural language (CNL). An interesting variation of this problem, and to the best of our knowledge, scarcely explored variation in this context, is when the input information is inconsistent. In such situations, the existing encodings of word puzzles produce inconsistent representations and break down. In this paper, we bring the well-known type of paraconsistent logics, called *Annotated Predicate Calculus* (APC) (Kifer and Lozinskii 1992), to bear on the problem. We introduce a new kind of non-monotonic semantics for APC, called *consistency preferred stable models* and argue that it makes APC into a suitable platform for dealing with inconsistency in word puzzles and, more generally, in NL sentences. We also devise a number of general principles to help the user choose among the different representations of NL sentences, which might seem equivalent but, in fact, behave differently when inconsistent information is taken into account. These principles can be incorporated into existing CNL translators, such as Attempto Controlled English (ACE) (Fuchs et al. 2008) and PENG Light (White and Schwitter 2009). Finally, we show that APC with the consistency preferred stable model semantics can be equivalently embedded in ASP with preferences over stable models, and we use this embedding to implement this version of APC in Clingo (Gebser et al. 2011) and its Asprin add-on (Brewka et al. 2015).

To appear in Theory and Practice of Logic Programming (TPLP).

KEYWORDS: Paraconsistency, Word Puzzles, Annotated Predicate Calculus, Controlled Natural Language

1 Introduction

The problem of logical representation for word puzzles has recently received considerable attention (Ponnuru et al. 2004; Shapiro 2011; Baral and Dzifcak 2012; Schwitter 2013). In all of these studies, however, the input information is assumed to be consistent and the proposed logical representations break on inconsistent input. The present paper proposes an approach that works in the presence of inconsistency and not just for word puzzles.

At first sight, one might think that the mere use of a paraconsistent logic such as Belanp's four valued logic (Belnap Jr 1977) or Annotated Logic Programming (Blair and Subrahmanian 1989; Kifer and Subrahmanian 1992) is all what is needed to address the problem, but it is not so. We do start with a well-known paraconsistent logic, called *Annotated Predicate Calculus* (APC) (Kifer and Lozinskii 1992), which is related to the aforementioned Annotated Logic Programs, but this is not enough: a number of issues arise in the presence of paraconsistency and different translations might seem equivalent but behave differently when inconsistent information is taken into account. As it turns out, several factors can affect the choice of the "right" logical representation for many NL sentences, especially for implications. We formalize several principles to

guide the translation of NL sentences into APC, principles that can be incorporated into existing controlled language translators, such as Attempto Controlled English (ACE) (Fuchs et al. 2008) and PENG Light (White and Schwitter 2009). We illustrate these issues with the classical Jobs Puzzle (Wos et al. 1984) and show how inconsistent information affects the conclusions.

To address the above problems formally, we introduce a new kind of non-monotonic semantics for APC, which is based on *consistency-preferred stable models* and is inspired by the concept of the most epistemically-consistent models of (Kifer and Lozinskii 1992). We argue that this new semantics makes APC into a good platform for dealing with inconsistency in word puzzles and, more generally, for translating natural language sentences into logic.

Finally, we show that the consistency-preferred stable models of APC can be computed using answer-set programming (ASP) systems that support preferences over stable models, such as Clingo (Gebser et al. 2011) with the Asprin add-on (Brewka et al. 2015).

This paper is organized as follows. Section 2 provides background material on APC. In Section 3 we consider the logic programming subset of APC and define preferential stable models for it. In Section 4, we show that the logic programming subset of APC (under the consistency-preferred stable model semantics) can be encoded in ASP in semantically-preserving way. In Section 5, we discuss variations of Jobs Puzzle (Wos et al. 1984) when various kinds of inconsistency are injected into the formulation of the puzzle. Section 6 explains that logical encoding of common knowledge in the presence of inconsistency needs to take into account a number of considerations that are not present when inconsistency is not an issue. We organize those considerations into several different principles and illustrate their impact. Section 8 concludes the paper. Finally, Appendix A contains the full encoding of Jobs Puzzle in APC under the consistency-preferred semantics. This appendix also includes variations that inject various kinds of inconsistency into the puzzle, and the derived conclusions are discussed. Appendices B and C contain similar analyses of other well-known puzzles: Zebra Puzzle[1] and Marathon Puzzle (C. Guéret and Sevaux 2000). Ready-to-run encodings of these programs in Clingo/Asprin can be found at `https://bitbucket.org/tiantiangao/apc_lp`.

2 Annotated Predicate Calculus: Background and Extensions

To make this paper self-contained, this section provides the necessary background on APC. At the end of the section, we define new semantic concepts for APC, which will be employed in later sections.

The *alphabet* of APC consists of countably-infinite sets of: *variables* $\mathscr{V}$, *function symbols* $\mathscr{F}$ (each symbol having an arity; *constants* are viewed as 0-ary function symbols), *predicate symbols* $\mathscr{P}$, *truth annotations*, quantifiers, and logical connectives. In (Kifer and Lozinskii 1992), truth annotations could come from an arbitrary upper semilattice (called "belief semilattice" there), but here we will use only $\bot$ (unknown), **f** (false), **t** (true) and $\top$ (contradiction or inconsistency), which are partially ordered as follows: $\bot \leq \mathbf{f} \leq \top$ and $\bot \leq \mathbf{t} \leq \top$. *Terms* in APC are constructed exactly as in predicate calculus: from constants, variables and function symbols. A *ground term* is one that has no variables.

Definition 1 (**Atomic formulas (Kifer and Lozinskii 1992)).** A *predicate term* has the form $p(t_1, t_2, \ldots, t_n)$, where p is a n-ary predicate symbol and t_1, t_2, …, t_n are terms. An APC *atomic*

[1] `https://en.wikipedia.org/wiki/Zebra_Puzzle`

formula (or an APC predicate) has the form $p(t_1,t_2,\ldots,t_n) : \mathbf{s}$, where $p(t_1,t_2,\ldots,t_n)$ is a predicate term and $\mathbf{s}$ is annotation indicating the degree of belief (or truth) in the predicate term. A *ground atomic formula* is an atomic formula that has no variables. □

We call an atomic formula of the form $p : \mathbf{s}$ a **t**-*predicate* (resp., an **f**-, $\top$-, or $\bot$-predicate) if $\mathbf{s}$ is **t** (resp., **f**-, $\top$-, or $\bot$).

APC includes the usual universal and existential quantifiers, the connectives, $\wedge$ and $\vee$, and there are two negation and two implication connectives: the *ontological negation* $\neg$ and *ontological implication* $\leftarrow$, plus the *epistemic negation* $\sim$ and *epistemic implication* $<\!\sim$. As will be seen later, the distinction between the ontological and the epistemic connectives is useful because they behave differently in the presence of inconsistency.

Definition 2 (**APC well-formed formulas (Kifer and Lozinskii 1992)).** An APC *well-formed formula* is defined inductively as follows:

- an atomic formula $p(t_1,t_2,\ldots,t_n) : \mathbf{s}$
- if ϕ and ψ are well-formed formulas, then so are $\sim\phi$, $\neg\phi$, $\phi\wedge\psi$, $\phi\vee\psi$, $\phi\leftarrow\psi$, and $\phi <\!\sim \psi$.
- if ϕ is a formula and X is a variable, then $(\forall X\phi)$ and $(\exists X\phi)$ are formulas. □

An APC *literal* is either a predicate $p : \mathbf{s}$ or an ontologically negated predicate $\neg p : \mathbf{s}$. An *epistemic literal* is either a predicate $p : \mathbf{s}$ or an epistemically negated predicate $\sim p : \mathbf{s}$.

In (Kifer and Lozinskii 1992), the semantics was defined with respect to general models, but here we will be dealing with logic programs and the Herbrand semantics will be more handy.

Definition 3 (**APC Herbrand universe, base, and interpretations).** The *Herbrand universe* $\mathscr{U}$ for APC is the set of all ground terms. The *Herbrand base* $\mathscr{B}$ for APC is the set of all ground APC atomic formulas. An *Herbrand interpretation* I for APC is a non-empty subset of the Herbrand base that is *closed* with respect to the following operations:

- if $p : \mathbf{s} \in I$, then also $p : \mathbf{s}' \in I$ for all $\mathbf{s}' \leq \mathbf{s}$; and
- if $p : \mathbf{s}_1$, $p : \mathbf{s}_2 \in I$, and $\mathbf{s} = lub_{\leq}(\mathbf{s}_1,\mathbf{s}_2)$ then $p : \mathbf{s} \in I$.

The annotations used in APC form a lattice (in our case a 4-element lattice) with the order $\leq$ and with $lub_{\leq}$ used as the least upper bound operator of that lattice.
We will also use $\mathscr{B}_{\top}$ to denote the subset of all $\top$-*predicates* in $\mathscr{B}$. □

As usual, a *variable assignment* is a mapping $\nu : \mathscr{V} \rightarrow \mathscr{U}$ that takes a variable and returns a ground term. This mapping is extended to terms as follows: $\nu(f(t_1,\ldots,t_n)) = f(\nu(t_1),\ldots,\nu(t_n))$. We will disregard variable assignments for formulas with no free variables (called *sentences*) since they do not affect ground formulas.

Definition 4 (**APC Herbrand Models).** Let I be an APC Herbrand interpretation and ν be a variable assignment. For an atomic formula $p(t_1,t_2,\ldots,t_n) : \mathbf{s}$, we write $I \models_\nu p(t_1,t_2,\ldots,t_n) : \mathbf{s}$ if and only if $p(\nu(t_1),\nu(t_2),\ldots,\nu(t_n)) : \mathbf{s} \in I$. For well-formed formulas ϕ and ψ, we write:

- $I \models_\nu \phi\wedge\psi$ if and only if $I \models_\nu \phi$ and $I \models_\nu \psi$;
- $I \models_\nu \phi\vee\psi$ if and only if $I \models_\nu \phi$ or $I \models_\nu \psi$;
- $I \models_\nu \neg\phi$ if and only if not $I \models_\nu \phi$;
- $I \models_\nu (\forall X)\phi$ if and only if $I \models_{\nu'} \phi$, for every assignment ν' that differs from ν only in its X-value;

- $I \models_\nu (\exists X)\phi$ if and only if $I \models_{\nu'} \phi$, for some ν' that differs from ν only in its X-value;
- $I \models_\nu \psi \leftarrow \phi$ if and only if $I \models_\nu \neg\phi \vee \psi$;
- $I \models_\nu \sim p : \mathbf{s}$ if and only if $I \models_\nu p : \sim \mathbf{s}$, where $\sim \mathbf{t} = \mathbf{f}$, $\sim \mathbf{f} = \mathbf{t}$, $\sim \top = \top$ and $\sim \bot = \bot$;

We also define: $\sim \neg\phi \equiv \neg \sim \phi$, $\sim (\phi \wedge \psi) \equiv \sim \phi \vee \sim \psi$, $\sim (\phi \vee \psi) \equiv \sim \phi \wedge \sim \psi$, $\sim \forall X\phi \equiv \exists X \sim \phi$, $\sim \exists X\phi \equiv \forall X \sim \phi$, and $\psi \mathrel{<\!\!\sim} \phi \equiv \sim \phi \vee \psi$.

A formula ϕ is *satisfied* by I if and only if $I \models_\nu \phi$ for every valuation ν. In this case we write simply $I \models \phi$. I is a *model* of a set of formulas P if and only if every formula $\phi \in P$ is satisfied in I. A set of formulas P logically entails a formula ψ, denoted $P \models \psi$, if and only if every model of P is also a model of ψ. □

APC has two types of logical entailment: *ontological* and *epistemic*. Ontological entailment is the entailment $\models$, which we have just defined. Before defining the *epistemic* entailment, we motivate it with a number of examples. To avoid clutter, in all examples we will only show the highest annotation for each APC predicate. For instance, if a model contains $p : \top$, then we will not show $p : \mathbf{t}$, $p : \mathbf{f}$, or $p : \bot$.

Example 1. *Consider the following set of APC formulas* $P = \{q : \mathbf{t} \leftarrow p : \mathbf{t},\ \ p : \mathbf{t}\}$. *It has four models:* $m_1 = \{p : \mathbf{t},\ q : \mathbf{t}\}$, $m_2 = \{p : \top,\ q : \mathbf{t}\}$, $m_3 = \{p : \mathbf{t},\ q : \top\}$ *and* $m_4 = \{p : \top,\ q : \top\}$. *Thus,* $P \models q : \mathbf{t}$ *holds (since* $q : \mathbf{t}$ *occurs in every model of* P*).* □

Example 2. *The APC set of formulas* $P = \{q : \mathbf{t} \leftarrow p : \mathbf{t},\ \ p : \top\}$ *has two models:* $m_1 = \{p : \top,\ q : \mathbf{t}\}$ *and* $m_2 = \{p : \top,\ q : \top\}$. *Therefore,* $P \models q : \mathbf{t}$ *holds.* □

Example 3. *This set of formulas* $P = \{q : \mathbf{t} \mathrel{<\!\!\sim} p : \mathbf{t},\ \ p : \mathbf{t}\}$ *is similar to that in Example 1 except that it uses epistemic implication instead of the ontological one. One of the models of that set is* $m = \{p : \top,\ q : \bot\}$ *and therefore* $P \not\models q : \mathbf{t}$. □

Examples 1 and 2 show that *ontological implication* has the modus ponens property, but it may be too strong, as it allows one to draw conclusions from inconsistent information. *Epistemic implication* of Example 3, on the other hand, is too cautious and does *not* have the modus ponens property. However, epistemic implication *does* have the modus ponens property *and* it blocks drawing conclusions from inconsistency under the *epistemic entailment*, defined next.

Definition 5 (**Most *e*-consistent models (Kifer and Lozinskii 1992)).** A Herbrand interpretation I_1 is *more* (or equally) *e-consistent* than another interpretation I_2 (denoted $I_1 \leq_\top I_2$) if and only if $I_1 \models p : \top$ implies $I_2 \models p : \top$ for every ground predicate term p.

A model I of a set of formulas P is a *most e-consistent model*, if there is no other model of P that is strictly more e-consistent than I.

A program P *epistemically* entails a formula ψ, denoted $P|\!\approx \psi$, if and only if every most e-consistent model of P is also a model of ψ. □

Going back to Example 3, it has only one most e-consistent model $m = \{p : \mathbf{t}, q : \mathbf{t}\}$, so $P|\!\approx q : \mathbf{t}$ holds. The next example shows that $\mathrel{<\!\!\sim}$ does *not* propagate inconsistency to conclusions.

Example 4. *Let* $P = \{q : \mathbf{t} \mathrel{<\!\!\sim} p : \mathbf{t},\ \ p : \top\}$. *Observe that* P *has a most e-consistent model* $m = \{p : \top,\ q : \bot\}$, *in which* $q : \mathbf{t}$ *does not hold. Therefore,* $P \not|\!\approx q : \mathbf{t}$ *holds.* □

Next we observe that not all inconsistent information is created equal, as people have different degrees of confidence in different pieces of information. For instance, one normally would have higher confidence in the fact that someone named Robin is a person than in the fact that Robin is

a male. Therefore, given a choice, we would hold it less likely that $person(robin)$ is inconsistent than that $male(robin)$ is. Likewise, in the following example, given a choice, we are more likely to hold to a belief that Pete is a person than to a belief that he is rich.

Example 5. *Consider the following formulas*

$$person(pete) : \mathbf{t}.$$
$$businessman(pete) : \mathbf{t}.$$
$$rich(pete) : \mathbf{f}.$$
$$rich(X) : \mathbf{t} \leftsquigarrow person(X) : \mathbf{t} \wedge businessman(X) : \mathbf{t}.$$

There are three most e-consistent models:

$$m_1 = \{person(pete) : \top,\ businessman(pete) : \mathbf{t},\ rich(pete) : \mathbf{f}\}$$
$$m_2 = \{person(pete) : \mathbf{t},\ businessman(pete) : \top,\ rich(pete) : \mathbf{f}\}$$
$$m_3 = \{person(pete) : \mathbf{t},\ businessman(pete) : \mathbf{t},\ rich(pete) : \top\}$$

Based on the aforesaid confidence considerations, we are more likely to believe that Pete is a person than that he is a businessman or rich. Therefore, we are likely to think that the models m_2 and m_3 are better descriptions of the real world than m_1. □

In this paper, we capture the above intuition by extending the notion of most e-consistent models with additional preferences over models.

Definition 6 (Consistency-preference relation and consistency-preferred models). A *consistency preference* S over interpretations, where S is a set of ground $\top$-predicates in APC, is defined as follows:

- An interpretation I_1 is *consistency-preferred* over I_2 with respect to S, denoted $I_1 <_S I_2$, if and only if $S \cap I_1 \subset S \cap I_2$.
- Interpretation I_1 and I_2 are *consistency-equal* with respect to S, denoted $I_1 =_S I_2$, if and only if $S \cap I_1 = S \cap I_2$.

A *consistency-preference relation* $<_{\mathscr{S}}$, where $\mathscr{S} = (S_1, S_2, \ldots, S_n)$ is a sequence of sets of ground $\top$-predicates, is defined as a lexicographic order composed out of the sequence of consistency preferences $S_1, S_2, \ldots, S_n$. Namely, $I_1 <_{\mathscr{S}} I_2$ if and only iff there is $1 \leq i \leq n$ such that $(\bigwedge_{1 \leq j < i}(I_1 =_{S_j} I_2))$ and $I_1 <_{S_i} I_2$.

A model I of a set of formulas P is called (most) *consistency-preferred* with respect to $<_{\mathscr{S}}$ if P has no other model J such that $J <_{\mathscr{S}} I$.

We will always assume that $S_n = \mathscr{B}_\top$ — the set of all ground $\top$-predicates and, therefore, any most consistency-preferred model is also a most e-consistent one.

We use the notation $|\approx_{\mathscr{S}}$ to denote *epistemic entailment* with respect to most consistency-preferred models. A program P *epistemically entails* a formula ψ with respect to a consistency-preference relation $<_{\mathscr{S}}$, denoted $P |\approx_{\mathscr{S}} \psi$, if and only if every most consistency-preferred model of P is also a model of ψ.

□

3 Logic Programming Subset of APC and Its Stable Models Semantics

In this section, we define the logic programming subset of APC, denoted APC_{LP}, and give it a new kind of semantics based on *consistency-preferred stable models*.

Definition 7. An *APC$_{LP}$ program* consists of rules of the form:

$$l_0 \vee \cdots \vee l_m \leftarrow l_{m+1} \wedge \cdots \wedge l_n \wedge \neg\, l_{n+1} \wedge \cdots \wedge \neg\, l_k.$$

where each l_i is an epistemic literal. Variables are assumed to be implicitly universally quantified. An *APC$_{LP}$ formula* is either a singleton epistemic literal, or a conjunction of epistemic literals, or a disjunction of them. □

The formula $l_0 \vee \cdots \vee l_m$ is called the *head* of the rule, and $l_{m+1} \wedge \cdots \wedge l_n \wedge \neg\, l_{n+1} \wedge \cdots \wedge \neg\, l_k$ is the *body* of that rule.

Recall from Section 2 that epistemic negation can be pushed inside and eliminated via this law: $\sim p : \alpha \equiv p :\sim \alpha$, where $\sim \mathbf{t} = \mathbf{f}$, $\sim \mathbf{f} = \mathbf{t}$, $\sim \top = \top$, and $\sim \bot = \bot$ so, for brevity, we assume that all *APC$_{LP}$* programs are transformed in this way and the epistemic negation is eliminated.

When the rule body is empty, the ontological implication symbol $\leftarrow$ is usually omitted and the rule becomes a disjunction. Such a disjunction can also be represented as an epistemic implication and sometimes this representation may be closer to a normal English sentence. For instance, the sentence, "If a person is a businessman then that person is rich," can be represented as an epistemic implication: $\mathit{rich}(X) : \mathbf{t} \;<\!\sim\; \mathit{person}(X) : \mathbf{t} \wedge \mathit{businessman}(X) : \mathbf{t}$, which is easier to read than the equivalent disjunction $\mathit{rich}(X) : \mathbf{t} \vee \mathit{person}(X) : \mathbf{f} \vee \mathit{businessman}(X) : \mathbf{f}$.

The notion of stable models for *APC$_{LP}$* carries over from standard answer set programming (ASP) with very few changes.

Definition 8 (The Gelfond-Lifschitz reduct for *APC$_{LP}$*). Let P be an *APC$_{LP}$* program and M be a Herbrand interpretation. The reduct of P w.r.t. M, denoted $\frac{P}{M}$, is a program free from ontological negation obtained by

1. removing rules with $\neg p : \mathbf{s}$ in the body, where $M \models p : \mathbf{s}$; and
2. removing literals $\neg p : \mathbf{s}$ from all remaining rules. □

Definition 9 (Stable models for *APC$_{LP}$*). A Herbrand interpretation M is a *stable model* of an *APC$_{LP}$* program P if M is a minimal model of $\frac{P}{M}$. Here, minimality is with respect to set inclusion. □

Definition 10 (Consistency-preferred stable models for *APC$_{LP}$*). Let $<_{\mathscr{S}}$ be a consistency-preference relation of Definition 6, where $\mathscr{S} = (S_1, S_2, \ldots, S_n)$ is a sequence of sets of ground $\top$-predicates. An *APC$_{LP}$* interpretation M is a (most) *consistency-preferred stable model* of an *APC$_{LP}$* program P if and only if:

- M is a stable model of P, and
- M is a most consistency-preferred model with respect to $<_{\mathscr{S}}$.

4 Embedding *APC$_{LP}$* into ASP

We now show that *APC$_{LP}$* can be isomorphically embedded in ASP extended with a model preference framework, such as the Clingo system (Gebser et al. 2011) with its Asprin extension (Brewka et al. 2015). We then prove the correctness of this embedding, i.e., that it is one-to-one and preserves the semantics. Next, we define the subset of ASP onto which *APC$_{LP}$* maps.

Definition 11. ASP_{truth} is a subset of *ASP* programs where the only predicate is `truth/2`, which is used to reify the APC predicate terms and associate them with truth values. That is, these atoms have the form $truth(p, \mathbf{s})$, where the first argument is the reification of an APC predicate term and the second argument is one of these truth annotations: `t`, `f`, `top`, or `bottom`.

An ASP_{truth} *program* consists a set of rules of the form:

$$t_0 \vee \cdots \vee t_m \leftarrow t_{m+1} \wedge \cdots \wedge t_n \wedge not\ t_{n+1} \wedge \cdots \wedge not\ t_k.$$

where the t_i's are `truth/2`-predicates.

An ASP_{truth} *formula* is either a singleton `truth/2`-predicate, a conjunction of such predicates, or a disjunction of them. □

Definition 12. The embedding of an APC_{LP} program in ASP_{truth}, denoted Ξ_{asp}, is defined recursively as follows (where tv_{asp} is the truth value mapping):

- $tv_{asp}(\mathbf{t}) = \texttt{t}$
- $tv_{asp}(\mathbf{f}) = \texttt{f}$
- $tv_{asp}(\top) = \texttt{top}$
- $tv_{asp}(\bot) = \texttt{bottom}$
- $\Xi_{asp}(p : \mathbf{s}) = \texttt{truth(p,}tv_{asp}(\mathbf{s})\texttt{)}$
- $\Xi_{asp}(\neg L) = \texttt{not}\ \Xi_{asp}(L)$, where L is an APC predicate
- $\Xi_{asp}(L \vee \phi) = \Xi_{asp}(L) \vee \Xi_{asp}(\phi)$, where L is an APC predicate and ϕ is a disjunction of APC predicates
- $\Xi_{asp}(L \wedge \phi) = \Xi_{asp}(L) \wedge \Xi_{asp}(\phi)$, where L is an APC literal and ϕ is a conjunction of APC literals
- $\Xi_{asp}(head \leftarrow body) = \Xi_{asp}(head) \leftarrow \Xi_{asp}(body)$, where *head* (resp., *body*) denotes the head (resp., the body) of a rule.

The embedding Ξ_{asp} also applies to APC Herbrand interpretations: each APC Herbrand interpretation (which is a set of APC atoms of the form $p : \mathbf{s}$) is mapped to a set of ASP_{truth} atoms (of the form `truth(p,`$tv_{asp}(\mathbf{s})$`)`). □

We require that each ASP_{truth} program includes the following background axioms to match the semantics of APC:

```
truth(X,top) :- truth(X,t),truth(X,f).
truth(X,t) :- truth(X,top).
truth(X,f) :- truth(X,top).
truth(X,bottom).
```

Lemma 1. *The embedding* $\Xi asp : APC_{LP} \rightarrow ASP_{truth}$ *is a one-to-one correspondence.* □

Proof. As mentioned, we can limit our attention to $\sim$-free programs. First, it is obvious that Ξ_{asp} is injective on APC literals. Injectivity on APC conjunctions and disjunctions can be shown by a straightforward induction on the number of conjuncts and disjuncts. Surjectivity follows similarly because it is straightforward to define the inverse of Ξ_{asp} by reversing the equations of Definition 12. □

Next, we show the above APC-to-ASP embedding preserves models, Gelfond-Lifshitz reduct, stable models, and also consistency preference relations.

Lemma 2. *The models of any ASP_{truth} program are closed with respect to $lub_{\leq}$ and downward-closed with respect to the $\leq$-ordering. Also, I is a model of an APC_{LP} program P if and only if $\Xi_{asp}(I)$ is a model of $\Xi_{asp}(P)$.*

Proof. Recall that every APC_{truth} is required to have the four rules listed right after Definition 12. These rules obviously enforce the requisite closures. The second part of the lemma follows directly from the definitions. □

Lemma 3. *Ξ_{asp} preserves the Gelfond-Lifshitz reduct: $\Xi_{asp}(\frac{P}{I}) = \frac{\Xi asp(P)}{\Xi_{asp}(I)}$.* □

Proof. For every predicate $p : \mathbf{s} \in P$, we have $I \models p : \mathbf{s}$ if and only if $\Xi_{asp}(I) \models \mathit{truth}(p, \mathbf{s})$, by Lemma 2. By the same lemma, if $r \in P$ then $I \models p : \mathbf{s}$ where $\neg p : \mathbf{s} \in \mathit{body}(r)$ if and only if $\Xi_{asp}(I) \models \mathit{truth}(p, \mathbf{s})$, where *not* $\mathit{truth}(p, \mathbf{s}) \in \mathit{body}(\Xi_{asp}(r))$. As a result, rule r gets eliminated by Gelfond-Lifschitz reduction if and only if $\Xi_{asp}(r)$ is eliminated and a negative literal in the body of r gets dropped if and only if its image in $\Xi_{asp}(r)$ gets dropped. □

Lemma 4. *Let I be a APC Herbrand interpretation. J is an APC Herbrand model of $\frac{P}{I}$ if and only if $\Xi_{asp}(J)$ is a model of $\Xi_{asp}(\frac{P}{I})$.* □

Proof. If $r \in P$ is a rule then $J \models \mathit{head}(r)$ if and only if $\Xi_{asp}(J) \models \mathit{head}(\Xi_{asp}(r))$ and $J \not\models \mathit{body}(r)$ if and only if $\Xi_{asp}(J) \not\models \mathit{body}(\Xi_{asp}(r))$. Thus, $J \models r$ if and only if $\Xi_{asp}(J) \models \Xi_{asp}(r)$. □

Lemma 5. *Let I_1 and I_2 be APC Herbrand interpretations. $I_1 \subseteq I_2$ if and only if $\Xi_{asp}(I_1) \subseteq \Xi_{asp}(I_2)$.* □

Proof. Follows directly from the definition of Ξ_{apc} and its inverse. □

Theorem 6. *M is a stable model of an APC_{LP} program P if and only if $\Xi_{asp}(M)$ is a stable model of $\Xi_{asp}(P)$.* □

Proof. By Lemma 4, J is a model of $\frac{P}{M}$ if and only if $\Xi_{asp}(J)$ is a model of $\frac{\Xi_{asp}(P)}{\Xi_{asp}(M)}$. Thus, the set of models for $\frac{P}{M}$ is in a one-one correspondence with the set of models for $\frac{\Xi_{asp}(P)}{\Xi_{asp}(M)}$. By Lemma 5, this correspondence preserves set-inclusion, so the set of *minimal* models of $\frac{P}{M}$ stands in one-one correspondence with respect to Ξ_{asp} with the set of minimal models of $\frac{\Xi_{asp}(P)}{\Xi_{asp}(M)}$. □

A consistency preference relation $<_{\mathscr{S}}$, where $\mathscr{S} = (S_1, S_2, \ldots, S_n)$, is translated into the following Asprin (Brewka et al. 2015) *lexico* preference relation $\mathscr{A}$ along with several `subset` preferences relations, each corresponding to one of the $<_{S_i}$ that are part of $<_{\mathscr{S}}$ (see Definition 6).

`#preference(`$\mathscr{A}$ `, lexico){ 1::name(`s_1`); 2::name(`s_2`); ...; n::name(`s_n`)}.`
`#preference(`s_1 `, subset){`the list of elements in $\Xi_{asp}(S_1)$`}.`
...
`#preference(`s_n `, subset){`the list of elements in $\Xi_{asp}(S_n)$`}.`

Lemma 7. *Let I_1 and I_2 be APC Herbrand interpretations, $\mathscr{S} = (S_1, S_2, \ldots, S_n)$ be a consistency preference relation and $\mathscr{A}$ be its corresponding Asprin preference relation. $I_1 <_{\mathscr{S}} I_2$ if and only if $\Xi_{asp}(I_1)$ is preferred over $\Xi_{asp}(I_2)$ with respect to $\mathscr{A}$.* □

Proof. The definition in the Asprin manual of the Asprin *lexico* and *subset* preference relations, as applied to our preference statements $\mathscr{A}$ given just prior to Lemma 7, is just a paraphrase of the lexicographical consistency-preference relation $\mathscr{S}$ in Definition 6. The lemma now follows from the obvious fact that Ξ_{asp} maps $\top$-literals of ASP onto the `top`-literals of ASP_{truth}, which have the form $\mathit{truth}(p, \mathtt{top})$. □

Theorem 8. *M is a most consistency preferred stable model of an APC$_{LP}$ program P with respect to a consistency preference relation $<_{\mathscr{S}}$ (where $\mathscr{S} = (S_1, S_2, \ldots, S_n)$) if and only if $\Xi_{asp}(M)$ is a preferred model of $\Xi_{asp}(P)$ with respect to the corresponding Asprin preference relation $\mathscr{A}$.* □

Proof. By Lemma 4, J is a model of $\frac{P}{M}$ if and only if $\Xi_{asp}(J)$ is a model of $\frac{\Xi_{asp}(P)}{\Xi_{asp}(M)}$. Since, by Lemma 7, Ξ_{asp} maps the preference relation $\mathscr{S}$ over the APC models into the preference relation $\mathscr{A}$ over the ASP models, the result follows. □

5 Jobs Puzzle and Inconsistency

Jobs Puzzle (Wos et al. 1984) is a classical logical puzzle that became a benchmark of sorts for many automatic theorem provers (Shapiro 2011; Schwitter 2013); it is also included in TPTP.[2] The usual description of Jobs Puzzle does not include implicit knowledge, like the facts that a person is either a male or a female (but not both), the husband of a person must be unique, etc., so we add this knowledge explicitly, like (Schwitter 2013). We also changed the name Steve to Robin in order to better illustrate one form of inconsistency.

1. There are four people: Roberta, Thelma, Robin and Pete.
2. Among them, they hold eight different jobs.
3. Each holds exactly two jobs.
4. The jobs are: chef, guard, nurse, telephone operator, police officer (gender not implied), teacher, actor, and boxer.
5. The job of nurse is held by a male.
6. The husband of the chef is the telephone operator.
7. Roberta is not a boxer.
8. Pete has no education past the ninth grade.
9. Roberta, the chef, and the police officer went golfing together.

In sum there are four people and eight jobs and to solve the puzzle one must figure out who holds which jobs. The solution is that Thelma is a chef and a boxer (and is married to Pete). Pete is a telephone operator and an actor. Roberta is a teacher and a guard. Finally, Robin is a police officer and a nurse.

However, if we inject inconsistency into the puzzle, current logical approaches fail because they are based on logics that do not tolerate inconsistency. Consider the following examples.

Example 6. *Let us add to the puzzle that "Thelma is an actor." Given that the original puzzle implies that Thelma is* not *an actor (she was a chef and a boxer), this addition causes inconsistencies. A first-order encoding of the puzzle (as, say, in TPTP) or an ASP-based one in (Schwitter 2013) will not find any models. In contrast, an encoding in APC$_{LP}$ can isolate inconsistent information. There are two possibilities: one where Thelma is an actor and the other where Thelma is a female. If we add background knowledge that Thelma is a female's name, it is less likely that Thelma's gender will be inconsistent, so the only consistency-preferred model will have one inconsistent conclusion that Thelma is an actor, but all other true facts will remain consistent.*

[2] Thousands of Problems for Theorem Provers (`http://www.cs.miami.edu/~tptp/`).

Example 7. *Consider adding the sentences "Robin is a male name" and "Robin is a female name," which will imply that Robin is both a male and female. The first-order and ASP-based encodings will, again, find no models, while an APC_{LP}-based encoding will localize inconsistency to just $male(robin) : \top$ and $female(robin) : \top$.*

Example 8. *Consider adding the sentence "Robin is Thelma's husband." Since the original job puzzle implies that Pete is Thelma's husband, this will cause inconsistency. If we add the background knowledge that husband is unique, again, the encoding of this modified puzzle in ACP_{LP} will localize inconsistency to just the aforesaid husband-facts.*

6 Knowledge Representation Principles for Inconsistency

Mere encoding of Jobs Puzzle in APC_{LP} is not enough because it is not unique: when inconsistency is taken into account, more information needs to be provided to obtain the encodings that match user intent. The main problem is that, if inconsistency is allowed, the number of possible worlds can grow to many hundreds even in relatively simple scenarios like Jobs Puzzle, and this practically annuls the benefits of the switch to a paraconsistent logic. We have already seen small examples of such scenarios at the end of Section 2, which motivated our notion of consistency preference, but there are more. We organize these scenarios around six main principles.

Principle 1: Contrapositive inference

Like in classical logic, contrapositive inference may be useful for knowledge representation. Consider the following sentences:

> If someone is a nurse, then that someone is educated.
> Pete is not educated.

We could encode the first sentence as $educated(X) : \mathbf{t} \leftarrow nurse(X) : \mathbf{t}$ or as $educated(X) : \mathbf{t} <\!\sim nurse(X) : \mathbf{t}$. Classically, the above sentences imply that Pete is not a nurse, but the encoding of the first sentence using the ontological implication $\leftarrow$ would not allow for that. If contrapositive inference is required, epistemic implication should be used.

Example 9. *Consider $P = \{educated(X) : \mathbf{t} <\!\sim nurse(X) : \mathbf{t}, \quad \sim educated(pete) : \mathbf{t}\}$. It has only one most consistency preferred model with respect to $<_{\mathscr{S}}$ (with $\mathscr{S} = (\mathscr{B}_\top)$), namely $m = \{educated(pete) : \mathbf{f}, \ nurse(pete) : \mathbf{f}\}$. Therefore, $P|\!\approx_{\mathscr{S}} nurse(pete) : \mathbf{f}$ holds.*

The above example uses contrapositive inference, but this is not always desirable. For instance, suppose $P' = \{male(X) : \mathbf{t} \leftarrow nurse(X) : \mathbf{t}, \quad \sim male(robin) : \mathbf{t}\}$. Here we use ontological implication to block contrapositive inference. Observe that P' has a most consistency preferred model with respect to $<_{\mathscr{S}}$, namely $m = \{male(robin) : \mathbf{f}, \ nurse(robin) : \bot\}$. Therefore, $nurse(robin) : \mathbf{f}$ does not hold, and this is exactly what we want, even if Robin happens to be not a male.[3] □

Principle 2: Propagation of inconsistency

As discussed in Example 2, APC gives us a choice of whether to draw conclusions from inconsistent information or not, and it is a useful choice. One way to block such inferences, illustrated in that example, is to use epistemic implication. Another way is to use the ontological implication with the $\mathbf{t} + \neg\top$ pattern in the rule body, e.g.,

$$educated(X) : \mathbf{t} \ \leftarrow \ nurse(X) : \mathbf{t} \ \wedge \ \neg nurse(X) : \top.$$

[3] In the USA as opposed to the U.K.

Both techniques block inferences from inconsistent information, but the second also blocks inference by contraposition, as discussed in Principle 1. The following examples illustrate the use of both of these methods.

Example 10. *Let* $P = \{educated(X) : \mathbf{t} <\!\sim nurse(X) : \mathbf{t},\ \ nurse(pete) : \top\}$. *Observe that there is one most consistency preferred model with respect to* $<_{\mathscr{S}}$ *(as before,* $\mathscr{S} = (\mathscr{B}_\top)$*):* $m = \{educated(pete) : \bot,\ \ nurse(pete) : \top\}$. *Therefore,* $P \not\approx_{\mathscr{S}} educated(pete) : \mathbf{t}$. □

Example 11. *Let* $P = \{educated(X) : \mathbf{t} \leftarrow nurse(X) : \mathbf{t} \wedge \neg nurse(X) : \top,\ \ nurse(pete) : \top\}$. *As in the previous example, P has a most consistency preferred model* $m = \{nurse(pete) : \top,\ \ educated(pete) : \bot\}$ *and so* $P \not\approx_{\mathscr{S}} educated(pete) : \mathbf{t}$. □

In both of these examples, inconsistency is not propagated through the rules, but Example 10 allows for contrapositive inference, while Example 11 does not. Indeed, suppose that instead of $nurse(pete) : \top$ we had $educated(pete) : \mathbf{f}$. Then, in the first case, $nurse(pete) : \mathbf{f}$ would be derived, while in the second it would not.

Blocking contrapositive inference and non-propagation of inconsistency can be applied selectively to some literals but not the others.

Example 12. *Consider the following sentence, "if a person holds a job of nurse then that person is educated". It can be encoded as*

$$(educated(X) : \mathbf{t} <\!\sim nurse(X) : \mathbf{t}) \leftarrow person(X) : \mathbf{t}.$$

The rule allows propagation of inconsistency through the person-predicate but blocks such propagation for the nurse-predicate. It also inhibits contrapositive inference of $person(pete) : \mathbf{f}$ *if the head of the rule is falsified by the additional facts* $nurse(pete) : \mathbf{t}$ *and* $educated(pete) : \mathbf{f}$. *However, due to the head of the rule, contrapositive inference would be allowed for* $nurse(pete) : \mathbf{f}$ *if* $educated(pete) : \mathbf{f}$ *was given.*

Principle 3: Polarity

This principle addresses situations such as the sentence "A person must be either a male or a female, but not both". When inconsistency is possible, we want to say three things: that any person must be either a male and or a female, that these facts cannot be unknown, and that if one of these is inconsistent then the other is too.

Example 13. *Let P be:*

$$\begin{array}{l} male(X) : \mathbf{t} \vee female(X) : \mathbf{t} \leftarrow person(X) : \mathbf{t} \\ male(X) : \mathbf{f} \vee female(X) : \mathbf{f} \leftarrow person(X) : \mathbf{t} \\ male(X) : \top \leftarrow person(X) : \mathbf{t} \wedge female(X) : \top \\ female(X) : \top \leftarrow person(X) : \mathbf{t} \wedge male(X) : \top \\ person(robin) : \mathbf{t} \end{array}$$

Two most consistency preferred models exist, which minimize the inconsistency of $person(robin)$*:*

$m_1 = \{person(robin) : \mathbf{t},\ \ male(robin) : \mathbf{t},\ \ female(robin) : \mathbf{f}\}$, and
$m_2 = \{person(robin) : \mathbf{t},\ \ male(robin) : \mathbf{f},\ \ female(robin) : \mathbf{t}\}$.

If we add $male(robin) : \top$ *(or* $female(robin) : \top$*) to P, then only one most consistency preferred model remains:* $m = \{male(robin) : \top,\ female(robin) : \top,\ person(robin) : \mathbf{t}\}$. □

Conditional *polarity* (or *polar dependency*) is generally represented as follows

$$p : \mathbf{t}\ ;\ q : \mathbf{t} \leftarrow \mathit{condition}.$$
$$p : \mathbf{f}\ ;\ q : \mathbf{f} \leftarrow \mathit{condition}.$$
$$q : \top \leftarrow \mathit{condition} \wedge p : \top.$$
$$p : \top \leftarrow \mathit{condition} \wedge q : \top.$$

where *condition* is a conjunction of atomic formulas and p, q are polar facts with respect to that condition.

Principle 4: Consistency preference relations

Recall from Example 5 that inconsistent information is not created equal, as people have different degrees of confidence in different pieces of information. For example, we have more confidence that someone whom we barely know is a person compared to the information about this person's marital situation (e.g., whether a husband exists). Therefore, person-facts are more likely to be consistent than marriage-facts and so we need to define consistency preference relations to specify the degrees of confidence. Consistency preference relations were introduced in Definition 6, and we already had numerous examples of its use. In Jobs Puzzle encoding in Appendix A, we use one, fairly elaborate, consistency preference relation. It first sets person and job information to be of the highest degree of confidence. Then, it prefers consistency of gender information of everybody but Robin. Third, it prefers consistency of the job assignment information. And finally, it minimizes inconsistency in general, for all facts.

Principle 5: Complete knowledge

This principle stipulates that certain information is defined completely, and cannot be unknown ($\bot$). But it can be inconsistent. Moreover, similarly to closed world assumption, negative information is preferred. For instance, if we do not know that someone is someone's husband, we may assume that that person is not. Such conclusions can be specified via a rule like this:

$$\mathit{husband}(X,Y) : \mathbf{f} \ \leftarrow\ \mathit{person}(X) : \mathbf{t} \ \wedge\ \mathit{person}(Y) : \mathbf{t} \ \wedge\ \neg \mathit{husband}(X,Y) : \mathbf{t}$$

Note that, unlike, say ASP, jumping to negative conclusions is not ensured by the stable model semantics of APC and must be given explicitly. But the advantage is that it can be done selectively. More generally, this type of reasoning can be specified as

$$p : \mathbf{f} \ \leftarrow\ \mathit{condition} \ \wedge\ \neg p : \mathbf{t}.$$

if p is known to be a predicate that is defined completely under the *condition*.

Principle 6: Exactly N

This principle captures the encoding of cardinality constraints in the presence of inconsistency. For instance, in Jobs Puzzle, the sentences "Every person holds exactly two jobs" and "Every job is held by exactly one person" are encoded as cardinality constraints:

$$2\ \{\mathit{hold}(X,Y) : \mathbf{t}\ \ \mathit{if}\ \ \mathit{job}(Y) : \mathbf{t}\}\ 2\ \leftarrow\ \mathit{person}(X) : \mathbf{t}.$$
$$1\ \{\mathit{hold}(X,Y) : \mathbf{t}\ \ \mathit{if}\ \ \mathit{person}(X) : \mathbf{t}\}\ 1\ \leftarrow\ \mathit{job}(Y) : \mathbf{t}.$$
$$\mathit{hold}(X,Y) : \mathbf{f}\ \leftarrow\ \mathit{person}(X) : \mathbf{t} \ \wedge\ \mathit{job}(Y) : \mathbf{t} \ \wedge\ \neg \mathit{hold}(X,Y) : \mathbf{t}.$$

These constraints count both true and inconsistent hold-facts, but can be easily modified to count only consistent true facts. Note the role of the last rule, which closes off the information being counted by the constraint. This is necessary because if, say, Pete is concluded to hold *exactly* two

jobs (of an actor and a phone operator) then there should be nothing unknown about him holding any other job. Instead, $hold(pete,X) : \mathbf{f} \wedge hold(pete,X) : \top$ should be true for any other job X.

The general form of the exactly N constraint is:

$$N\ \{L : \mathbf{t} \quad if \quad condition\}\ N \leftarrow guard.$$
$$L : \mathbf{f} \leftarrow guard \wedge condition \wedge \neg L : \mathbf{t}.$$

As in ASP, such statements can be represented as a number of ground disjunctive rules. The "exactly N" constraints can be generalized to "at least N and at most M" constraints, if we extend the semantics in the direction of (Soininen et al. 2001).

7 Comparison with Other Work

Although a great deal of work is dedicated to paraconsistent logics and logical formalizations for word puzzles separately, we are unaware of any work that applies paraconsistent logics to solving word puzzles that might contain inconsistencies. As we demonstrated, mere encoding of such puzzles in a paraconsistent logic leads to an explosion of possible worlds, which is not helpful.[4] Most paraconsistent logics (Priest et al. 2015; J. Y. Beziau 2007; Belnap Jr 1977; da Costa 1974) deal with inconsistency from the philosophical or mathematical point of view and do not discuss knowledge representation. Other paraconsistent logics (Blair and Subrahmanian 1989; Kifer and Subrahmanian 1992) were developed for definite logic programs and cannot be easily applied to solving more complex knowledge representation problems that arise in word puzzles. An interesting question is whether our use of APC is essential, i.e., whether the notions of consistency-preferred models can be adapted to other paraconsistent logics and the relationship with ASP can be established. First, it is clear that such an adaptation is unlikely for proof-theoretic approaches to inconsistency, such as (da Costa 1974). We do not know if such an adaptation is possible for model-theoretic approaches, such as (Belnap Jr 1977).

On the word puzzles front, (Wos et al. 1984) used the first-order logic theorem prover OTTER to solve Jobs Puzzle[5] and (Shapiro 2011) represented Jobs Puzzle in multiple logical languages: TPTP,[6] Constraint Lingo (Finkel et al. 2004) layered on top of the ASP system Smodels (Syrjänen and Niemelä 2001) as the backend, and the SNePS commonsense reasoning system (Shapiro 2000). More recently, (Baral and Dzifcak 2012; Schwitter 2013) represented word puzzles using NL/CNL sentences, and then automatically translate them into ASP. None of these underlying formalisms, FOL, ASP, and SNePS, are equipped to reason in the presence of inconsistency. In contrast, APC_{LP}, combined with the knowledge representation principles developed in Section 6, localizes inconsistency and computes useful possible worlds. In addition, APC_{LP} has mechanisms to control how inconsistency is propagated through inference, it allows one to prioritize inconsistent information, and it provides several other ways to express user's intent (through contraposition, completion of knowledge, etc.).

8 Conclusion

In this paper we discussed the problem of knowledge representation in the presence of inconsistent information with particular focus on representing English sentences using logic, as in word

[4] Also see Appendix A and the ready-to-run examples at `https://bitbucket.org/tiantiangao/apc_lp`.
[5] `http://www.mcs.anl.gov/~wos/mathproblems/jobs.txt`
[6] `http://www.cs.miami.edu/~tptp/cgi-bin/SeeTPTP?Category=Problems&Domain=PUZ&File=PUZ019-1.p`

puzzles (Wos et al. 1984; Shapiro 2011; Ponnuru et al. 2004; Schwitter 2013; Baral and Dzifcak 2012). We have shown that a number of considerations play a role in deciding on a particular encoding, which includes whether or not inconsistency should be propagated through implications, relative degrees of confidence in different pieces of information, and others. We used the well-known Jobs, Zebra and Marathon puzzles (see the appendices in the supplemental material) to illustrate many of the above issues and show how the conclusions change with the introduction of different kinds of inconsistency into the puzzle.

As a technical tool, we started with a paraconsistent logic called *Annotated Predicate Calculus* (Kifer and Lozinskii 1992) and then gave it a special kind of non-monotonic semantics that is based on *consistency-preferred stable models*. We also showed that these models can be computed using ASP systems that support preference relations over stable models, such as Clingo (Gebser et al. 2011) with the Asprin extension (Brewka et al. 2015).

For future work, we will consider additional puzzles which may suggest new knowledge representation principles. In addition, we will investigate ways to incorporate inconsistency into CNL systems. This will require introduction of background knowledge into these systems and linguistic cues into the grammar.

References

BARAL, C. AND DZIFCAK, J. 2012. Solving puzzles described in english by automated translation to answer set programming and learning how to do that translation. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012*. AAAI Press, Rome, Italy.

BELNAP JR, N. D. 1977. A useful four-valued logic. In *Modern uses of multiple-valued logic*. Springer, Volume 2, 5–37.

BLAIR, H. AND SUBRAHMANIAN, V. 1989. Paraconsistent logic programming. *Theoretical Computer Science 68*, 135–154.

BREWKA, G., DELGRANDE, J. P., ROMERO, J., AND SCHAUB, T. 2015. asprin: Customizing answer set preferences without a headache. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA*. AAAI Press, Austin, Texas, 1467–1474.

C. GUÉRET, C. P. AND SEVAUX, M. 2000. *Programmation linéaire - 65 problèmes d'optimisation modélisés et résolus avec Visual Xpress*. Eyrolles, France. ISBN : 2-212-09202-4.

DA COSTA, N. 1974. On the theory of inconsistent formal systems. *Notre Dame J. of Formal Logic 15*, 4 (October), 497–510.

FINKEL, R. A., MAREK, V. W., AND TRUSZCZYNSKI, M. 2004. Constraint lingo: towards high-level constraint programming. *Softw., Pract. Exper. 34*, 15, 1481–1504.

FUCHS, N. E., KALJURAND, K., AND KUHN, T. 2008. Attempto controlled english for knowledge representation. In *Reasoning Web, 4th International Summer School 2008, Venice, Italy, September 7-11, 2008, Tutorial Lectures*. Lecture Notes in Computer Science, vol. 5224. Springer, Venice, Italy, 104–124.

GEBSER, M., KAMINSKI, R., KAUFMANN, B., OSTROWSKI, M., SCHAUB, T., AND SCHNEIDER, M. 2011. Potassco: The Potsdam answer set solving collection. *AI Communications 24*, 2, 107–124.

J. Y. BEZIAU, W. CARNIELLI, D. M. G. 2007. *Handbook of Paraconsistency (Studies in Logic)*. College Publications, United States.

KIFER, M. AND LOZINSKII, E. L. 1992. A logic for reasoning with inconsistency. *J. Autom. Reasoning 9*, 2, 179–215.

KIFER, M. AND SUBRAHMANIAN, V. S. 1992. Theory of generalized annotated logic programming and its applications. *J. Log. Program. 12*, 3&4, 335–367.

PONNURU, H., FINKEL, R. A., MAREK, V. W., AND TRUSZCZYNSKI, M. 2004. Automatic generation of english-language steps in puzzle solving. In *Proceedings of the International Conference on Artificial*

Intelligence, IC-AI '04, June 21-24, 2004, Las Vegas, Nevada, USA, Volume 1. CSREA Press, Las Vegas, Nevada, USA, 437–442.

PRIEST, G., TANAKA, K., AND WEBER, Z. 2015. Paraconsistent logic. In *The Stanford Encyclopedia of Philosophy*, Spring 2015 ed., E. N. Zalta, Ed. Stanford, USA.

SCHWITTER, R. 2012. Answer set programming via controlled natural language processing. In *Controlled Natural Language - Third International Workshop, CNL 2012, August 29-31, 2012. Proceedings*, T. Kuhn and N. E. Fuchs, Eds. Lecture Notes in Computer Science, vol. 7427. Springer, Zurich, Switzerland, 26–43.

SCHWITTER, R. 2013. The Jobs Puzzle: Taking on the challenge via controlled natural language processing. *Theory and Practice of Logic Programming 13*, 4-5, 487–501.

SHAPIRO, S. C. 2000. An introduction to sneps 3. In *Conceptual Structures: Logical, Linguistic, and Computational Issues, 8th International Conference on Conceptual Structures, ICCS 2000, Darmstadt, Germany, August 14-18, 2000, Proceedings*. Springer, Darmstadt, Germany, 510–524.

SHAPIRO, S. C. 2011. The jobs puzzle: A challenge for logical expressibility and automated reasoning. In *Logical Formalizations of Commonsense Reasoning, Papers from the 2011 AAAI Spring Symposium, California, USA, March 21-23, 2011*. AAAI, Stanford, California, USA.

SOININEN, T., NIEMELÄ, I., TIIHONEN, J., AND SULONEN, R. 2001. Representing configuration knowledge with weight constraint rules. In *Answer Set Programming, Towards Efficient and Scalable Knowledge Representation and Reasoning, Proceedings of the 1st Intl. ASP'01 Workshop, Stanford, March 26-28, 2001*. Springer, Stanford, California, USA.

SYRJÄNEN, T. AND NIEMELÄ, I. 2001. The Smodels system. In *Logic Programming and Nonmonotonic Reasoning, 6th International Conference, LPNMR 2001, Vienna, Austria, September 17-19, 2001, Proceedings*. Springer, Vienna, Austria, 434–438.

WHITE, C. AND SCHWITTER, R. 2009. An update on PENG Light. In *Proceedings of ALTA*. Vol. 7. Springer, Sydney, Australia, 80–88.

WOS, L., OVERBECK, R., LUSK, E., AND BOYLE, J. 1984. *Automated reasoning: Introduction and applications*. Prentice Hall Inc.,Old Tappan, NJ, United States.

Appendix A Jobs Puzzle in APC_{LP} with Inconsistency Injections

We now present a complete APC_{LP} encoding of Jobs Puzzle and highlight the principles, introduced in Section 6, used in the encoding. We also show several cases of inconsistency injection and discuss the consequences. The English sentences are based on the CNL representation of Jobs Puzzle from Section 3 in (Schwitter 2013) where "Steve" is changed to "Robin" for the sake of an example (because Robin can be both a male and a female name).

1 Roberta is a person. Thelma is a person. Robin is a person. Pete is a person.
$person(roberta) : \mathbf{t}.\ \ person(thelma) : \mathbf{t}.\ \ person(robin) : \mathbf{t}.\ \ person(pete) : \mathbf{t}.$
2 Roberta is a female. Thelma is a female.
$female(roberta) : \mathbf{t}.\ \ female(thelma) : \mathbf{t}.$
3 Robin is male. Pete is male.
$male(robin) : \mathbf{t}.\ \ male(pete) : \mathbf{t}.$

Sentence 4 is encoded based on Principle 3, which treats *male* and *female* as polar facts.

4 Exclude that a person is male and that the person is female.
$male(X) : \mathbf{t} \vee female(X) : \mathbf{t} \leftarrow person(X) : \mathbf{t}.$
$male(X) : \mathbf{f} \vee female(X) : \mathbf{f} \leftarrow person(X) : \mathbf{t}.$
$female(X) : \top \leftarrow person(X) : \mathbf{t} \wedge male(X) : \top.$
$male(X) : \top \leftarrow person(X) : \mathbf{t} \wedge female(X) : \top.$

We encode sentences 5 and 6 using Principle 6, which constrains the cardinality of $hold(X, Y)$. This will cause second rule in Sentence 5 to be repeated as part of encoding of Sentence 6, so we omit the duplicate.

5 If there is a job then exactly one person holds that job.
$1\ \{hold(X,Y) : \mathbf{t} \quad if \quad person(X) : \mathbf{t}\}\ 1 \leftarrow job(Y) : \mathbf{t}.$
$hold(X,Y) : \mathbf{f} \leftarrow person(X) : \mathbf{t} \wedge job(Y) : \mathbf{t} \wedge \neg hold(X,Y) : \mathbf{t}.$
6 If there is a person then the person holds exactly two jobs.
$2\ \{hold(X,Y) : \mathbf{t} \quad if \quad job(Y) : \mathbf{t}\}\ 2 \leftarrow person(X) : \mathbf{t}.$

Encoding of the following facts is straightforward:

7 Chef is a job. Guard is a job. Nurse is a job. Operator is a job. Police is a job. Teacher is a job. Actor is a job. Boxer is a job.
$job(chef) : \mathbf{t}.\ \ job(guard) : \mathbf{t}.\ \ job(nurse) : \mathbf{t}.\ \ job(operator) : \mathbf{t}.$
$job(police) : \mathbf{t}.\ \ job(teacher) : \mathbf{t}.\ \ job(actor) : \mathbf{t}.\ \ job(boxer) : \mathbf{t}.$

Sentences 8-13 are encoded based on Principles 1 and 2, where contrapositive inference and propagation of inconsistency are allowed for some literals but not others. Notice that it is undesirable to allow propagation of inconsistency from *person*-facts and *job*-facts, since it is unreasonable to conclude that somebody is, say, a male while being unsure that this somebody is a person. Ditto about the jobs. Contrapositive reasoning (say, from non-male to non-person) is also inappropriate here because we have higher confidence in someone being a person. So, we use ontological implication $\leftarrow$ in the next group of rules.

8 If a person holds a job as a nurse then that person is a male.
$(male(X) : \mathbf{t} \ <\!\sim\ hold(X,nurse) : \mathbf{t}) \leftarrow person(X) : \mathbf{t} \wedge job(nurse) : \mathbf{t} \wedge \neg person(X) : \top \wedge \neg job(nurse) : \top.$

9 If a person holds a job as an actor then that person is a male.
$(male(X):\mathbf{t} \mathrel{<\!\!\sim} hold(X,actor):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge job(actor):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(actor):\top.$

10 If a first person holds a job as a chef and a second person holds a job as a telephone operator then the second person is a husband of the first person.
$(husband(Y,X):\mathbf{t} \mathrel{<\!\!\sim} hold(X,chef):\mathbf{t} \wedge hold(Y,operator):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge job(chef):\mathbf{t} \wedge person(Y):\mathbf{t} \wedge job(operator):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(chef):\top \wedge \neg person(Y): \wedge \neg job(operator):\top.$

11 If a first person is a husband of a second person then the first person is male.
$(male(X):\mathbf{t} \mathrel{<\!\!\sim} husband(X,Y):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge person(Y):\mathbf{t} \wedge \neg person(X):\top \wedge \neg person(Y):\top.$

12 If a first person is a husband of a second person then the second person is female.
$(female(Y):\mathbf{t} \mathrel{<\!\!\sim} husband(X,Y):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge person(Y):\mathbf{t} \wedge \neg person(X):\top \wedge \neg person(Y):\top.$

13 Exclude that Roberta holds a job as boxer.
$hold(roberta,boxer):\mathbf{f} \leftarrow job(boxer):\mathbf{t} \wedge \neg job(boxer):\top.$

Encoding of the following fact is straightforward.

14 Exclude that Pete is educated.
$educated(pete):\mathbf{f}.$

Sentences 15-20 are also encoded based on Principles 1 and 2.

15 If a person holds a job as nurse then the person is educated.
$(educated(X):\mathbf{t} \mathrel{<\!\!\sim} hold(X,nurse):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge job(nurse):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(nurse):\top.$

16 If a person holds a job as a police officer then that person is educated.
$(educated(X):\mathbf{t} \mathrel{<\!\!\sim} hold(X,police):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge job(police):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(police):\top.$

17 If a person holds a job as a teacher then the person is educated.
$(educated(X):\mathbf{t} \mathrel{<\!\!\sim} hold(X,teacher):\mathbf{t}) \leftarrow person(X):\mathbf{t} \wedge job(teacher):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(teacher):\top.$

18 Exclude that Roberta holds a job as a chef.
$hold(roberta,chef):\mathbf{f} \leftarrow job(chef):\mathbf{t} \wedge \neg job(chef):\top.$

19 Exclude that Roberta holds a job as a police officer.
$hold(roberta,police):\mathbf{f} \leftarrow job(police):\mathbf{t} \wedge \neg job(police):\top.$

20 Exclude that a person holds a job as a chef and that the same person holds a job as a police officer.
$hold(X,chef):\mathbf{f} \vee hold(X,police):\mathbf{f} \leftarrow person(X):\mathbf{t} \wedge job(chef):\mathbf{t} \wedge job(police):\mathbf{t} \wedge \neg person(X):\top \wedge \neg job(chef):\top \wedge \neg job(police):\top.$

Next we define the consistency preference relation $<_{\mathscr{S}}$, where $\mathscr{S} = (s_1, s_2, s_3, \mathscr{B}_\top)$, which implements Principle 4. Here s_1 says that we hold greater confidence in the information about someone being a person and something being a job than in any other kind of information in the puzzle. That is, these facts are least likely to be inconsistent. Next, s_2 says that we are very likely to believe that Pete is a male name and Thelma and Roberta are female names. We are not sure about Robin, so s/he is left out in s_2. The set s_3 says that next we are likely to believe the

information on who holds which jobs. The last component in $\mathscr{S}$, $\mathscr{B}_\top$, is the usual default that prefers the most e-consistent models.

$$s_1 = \{person(roberta) : \top,\ person(thelma) : \top,\ person(robin) : \top,\ person(pete) : \top,\ job(chef) : \top, job(guard) : \top, job(nurse) : \top, job(operator) : \top,\ job(police) : \top, job(teacher) : \top, job(actor) : \top, job(boxer) : \top\}$$

$$s_2 = \{male(pete) : \top,\ female(thelma) : \top,\ female(roberta) : \top\}$$
$$s_3 = \{hold(pete,chef) : \top,\ hold(pete,guard) : \top,\ hold(pete,nurse) : \top,\ hold(pete, operator) : \top,\ hold(pete,police) : \top,\ hold(pete,teacher) : \top,\ hold(pete, actor) : \top,\ hold(pete,boxer) : \top,\ \ldots\ \}$$

It is interesting to note that if s_3 is not included then, in some cases, there might be too many possibilities to solve the puzzle by allowing *hold*-predicates to be inconsistent. This is because many rules and constraints in the puzzle use *hold* as a premise. So, without minimizing the possibility of inconsistency in *hold* those rules and constraints become vacuously true, leading to an explosion of the possible worlds.

The APC_{LP} encoding generates one most consistency preferred model where the information concerning $hold/2$ is

$$hold(pete,actor) : \mathbf{t} \quad hold(pete,operator) : \mathbf{t} \quad hold(robin,police) : \mathbf{t}$$
$$hold(robin,nurse) : \mathbf{t} \quad hold(thelma,chef) : \mathbf{t} \quad hold(thelma,boxer) : \mathbf{t}$$
$$hold(roberta,guard) : \mathbf{t} \quad hold(roberta,teacher) : \mathbf{t}$$

Not surprisingly, this solution matches other approaches because so far we have not injected inconsistency (and so, for example, the preferences s_1, s_2, and s_3 do not matter here).

Next, we illustrate five cases of injection of inconsistency into the puzzle. Since complete models tend to be rather large, we show only $hold/2$ and $\top$-predicates.

21 Thelma is an actor — Variation 1
$hold(thelma,actor) : \mathbf{t}$.

This makes Thelma's job assignment as an actor inconsistent and we get this model:

$$m = \{hold(pete,boxer) : \mathbf{t} \quad hold(pete,operator) : \mathbf{t} \quad hold(robin,police) : \mathbf{t}$$
$$hold(robin,nurse) : \mathbf{t} \quad hold(thelma,actor) : \top \quad hold(thelma,chef) : \mathbf{t}$$
$$hold(roberta,teacher) : \mathbf{t} \quad hold(roberta,guard) : \mathbf{t}\}$$

Indeed, the original puzzle implies that Thelma is *not* an actor. Given that $hold(thelma,actor) : \mathbf{t}$ is true, Sentence 9 will sanction two possibilities: one where $male(thelma) : \mathbf{t}$ is true and the other where $hold(thelma,actor) : \top$ is true. In the first case, Sentence 4 will force $male(thelma) : \top$ and $female(thelma) : \top$ into the model. In the second case, we will have $hold(thelma,actor) : \top$ in the model. Given that we have high confidence in Thelma's gender (preference s_2), we hold her gender less likely to be inconsistent. Thus, the first case gets eliminated.

The next variation assumes that Robin is a female (instead of Thelma being an actor in Variation 1).

22 Robin is female — Variation 2
$female(robin) : \mathbf{t}$.

There are two models:

$m_1 = \{hold(pete,actor) : \mathbf{t}\ \ hold(pete,operator) : \mathbf{t}\ \ hold(robin,police) : \mathbf{t}$
$hold(robin,nurse) : \mathbf{t}\ \ hold(roberta,teacher) : \mathbf{t}\ \ hold(thelma,boxer) : \mathbf{t}$
$hold(roberta,guard) : \mathbf{t}\ \ hold(thelma,chef) : \mathbf{t}\ \ husband(pete,thelma) : \mathbf{t}$
$male(robin) : \top\ \ female(robin) : \top\}$

$m_2 = \{hold(pete,actor) : \mathbf{t}\ \ hold(pete,operator) : \mathbf{t}\ \ hold(robin,nurse) : \mathbf{t}$
$hold(robin,chef) : \mathbf{t}\ \ hold(thelma,boxer) : \mathbf{t}\ \ hold(thelma,police) : \mathbf{t}$
$hold(roberta,guard) : \mathbf{t}\ \ hold(roberta,teacher) : \mathbf{t}\ \ husband(pete,robin) : \mathbf{t}$
$female(robin) : \top\ \ male(robin) : \top\}$

Sentences 3 and 4 imply inconsistency regarding Robin's gender. The first model is the same as in the original puzzle (as far as the job assignments go). In the second model, since Robin's gender is inconsistent (both male and female), it is compatible to make Robin a chef and Thelma a police officer. Therefore, we derive that Pete is Robin's husband instead of Thelma's.

The next variation of the original puzzle explicitly assumes that Robin is Thelma's husband.

23 Robin is a husband of Thelma — Variation 3
$husband(robin,thelma) : \mathbf{t}$.

Here we need to add more background knowledge about marital relations. For instance, that every person can marry or be married to at most one person. Together with Sentence 23, this will cause inconsistency because the original puzzle implies that Pete is Thelma's husband. In (Schwitter 2013), this implicit knowledge is not stated, so it will fail to detect inconsistency. If such background knowledge were added to Schwitter's formulation as constraints then there would be no models. The background knowledge we need is:

24 A person who is a male is a husband of exactly one other person, or that person is null.
25 A person who is a female has exactly one husband or that husband is null.
26 Exclude that person X is a husband of Y and person Z is a husband of X simultaneously.
27 If it is not derivable that person X is person Y's husband, then X is not Y's husband.

Sentences 24 – 26 are cardinality constraints and are encoded based on Principle 6. Sentence 27 says that the information about husbands is complete; it is encoded based on Principle 5. Also, we block propagation of inconsistency from male and female based on Principle 2.

$1\ \{husband(X,Y) : \mathbf{t} \quad if \quad (person(Y) : \mathbf{t}\ or\ Y = null)\}\ 1\ \leftarrow$
$person(X) : \mathbf{t} \wedge male(X) : \mathbf{t} \wedge \neg male(X) : \top.$
$1\ \{husband(X,Y) : \mathbf{t} \quad if \quad (person(X) : \mathbf{t}\ or\ X = null)\}\ 1\ \leftarrow$
$person(Y) : \mathbf{t} \wedge female(Y) : \mathbf{t} \wedge \neg female(Y) : \top.$
$:\!-\ husband(X,Y) : \mathbf{t} \wedge husband(Z,X) : \mathbf{t} \wedge X\ !=\ null.$
$husband(X,Y) : \mathbf{f} \leftarrow person(X) : \mathbf{t} \wedge person(Y) : \ \wedge \neg husband(X,Y) : \mathbf{t}.$

There are six models. When these models are projected on $hold/2$ (which constitutes the solution to the puzzle) and the $\top$-predicates, we get three distinct sets:

$m_1 = \{hold(pete,actor) : \mathbf{t}\quad hold(pete,police) : \mathbf{t}\quad hold(robin,operator) : \mathbf{t}$
$hold(robin,nurse) : \mathbf{t}\quad hold(roberta,teacher) : \mathbf{t}\ \ hold(roberta,guard) : \mathbf{t}$
$hold(thelma,boxer) : \mathbf{t}\quad hold(thelma,chef) : \mathbf{t}\quad educated(pete) : \top\}$

$m_2 = \{hold(pete,actor) : \mathbf{t}\quad hold(pete,nurse) : \mathbf{t}\quad hold(robin,operator) : \mathbf{t}$
$hold(robin,police) : \mathbf{t}\quad hold(roberta,teacher) : \mathbf{t}\ \ hold(roberta,guard) : \mathbf{t}$

$$hold(thelma, boxer) : \mathbf{t} \quad hold(thelma, chef) : \mathbf{t} \quad educated(pete) : \top\}$$

$$\begin{aligned} m_3 = \{&hold(pete, police) : \mathbf{t} \quad hold(pete, nurse) : \mathbf{t} \quad hold(robin, operator) : \mathbf{t} \\ &hold(robin, actor) : \mathbf{t} \quad hold(roberta, teacher) : \mathbf{t} \quad hold(roberta, guard) : \mathbf{t} \\ &hold(thelma, boxer) : \mathbf{t} \quad hold(thelma, chef) : \mathbf{t} \quad educated(pete) : \top\} \end{aligned}$$

The puzzle originally implied that Pete is Thelma's husband. Since we now explicitly stated that Robin is Thelma's husband, Sentences 25 and 27 will force $husband(pete, thelma) : \mathbf{f}$ to hold. By Sentence 10, $hold(pete, operator) : \mathbf{t}$ and $hold(Thelma, chef) : \mathbf{t}$ cannot hold simultaneously, so many solutions with inconsistencies in them will be generated. Due to the consistency preference relations, the APC_{LP} encoding will prefer the models where $educated(pete) : \top$ holds.

Sentences 24 and 25 sanction two possibilities for the husband information in each of the above models. For instance, the model m_1 corresponds to two models out of the six models that we get; they differ only in their husband information. Given that Pete is not an operator and Roberta is not a chef, Pete is not necessarily Roberta's husband. Therefore, there are two cases: one where $husband(pete, roberta) : \mathbf{t}$ holds and the other where $husband(pete, null) : \mathbf{t}$ and $husband(null, roberta) : \mathbf{t}$ hold. Similar considerations apply to m_2 and m_3.

The next variation applies the background knowledge about husbands from Variation 3 to Variations 1 and 2.

28 Thelma is an actor. — Variation 4.1 (modification of Variation 1)
$hold(thelma, actor) : \mathbf{t}$.
Sentences 24 – 27.

There are now two models, and their projections on *hold*- and $\top$-predicates are compatible with the solution to Variation 1. The only difference between these two models is in *husband*-predicates, so we show only that part.

$$\begin{aligned} m_1 &= \{husband(pete, thelma) : \mathbf{t}, \quad husband(robin, roberta) : \mathbf{t}\} \\ m_2 &= \{husband(pete, thelma) : \mathbf{t}, \quad husband(robin, null) : \mathbf{t}, \quad husband(null, roberta) : \mathbf{t}\} \end{aligned}$$

29 Robin is a female. — Variation 4.2 (modification of Variation 2)
$female(robin) : \mathbf{t}$.
Sentences 24 – 27.

There are three models and their projections on *hold*- and $\top$-predicates are compatible with Variation 2. The only difference with Variation 2 is in the *husband*-predicates, which we show:

$$\begin{aligned} m_1 &= \{husband(pete, thelma) : \mathbf{t}, \quad husband(robin, roberta) : \mathbf{t}\} \\ m_2 &= \{husband(pete, thelma) : \mathbf{t}, \quad husband(robin, null) : \mathbf{t}, \quad husband(null, roberta) : \mathbf{t}\} \\ m_3 &= \{husband(pete, robin) : \mathbf{t}, \quad husband(null, thelma) : \mathbf{t}, \quad husband(null, roberta) : \mathbf{t}\} \end{aligned}$$

Here m_1 and m_2 correspond to the first model in Variation 2 and m_3 corresponds to the second model there.

Appendix B Zebra Puzzle in APC_{LP} with Inconsistency Injections

We now present a complete APC_{LP} encoding of the original Zebra Puzzle as is described in Wikipedia.[7] A slightly different version of the puzzle appears in TPTP.[8] The encoding highlights the principles introduced in Section 6 and we also discuss several cases of inconsistency injection.

(a) There are five houses.
(b) The Englishman lives in the red house.
(c) The Spaniard owns the dog.
(d) Coffee is drunk in the green house.
(e) The Ukrainian drinks tea.
(f) The green house is immediately to the right of the ivory house.
(g) The Old Gold smoker owns snails.
(h) Kools are smoked in the yellow house.
(i) Milk is drunk in the middle house.
(j) The Norwegian lives in the first house.
(k) The man who smokes Chesterfields lives in the house next to the man with the fox.
(l) Kools are smoked in the house next to the house where the horse is kept.
(m) The Lucky Strike smoker drinks orange juice.
(n) The Japanese smokes Parliaments.
(o) The Norwegian lives next to the blue house.

Now, who drinks water? Who owns the zebra?

In the interest of clarity, it must be added that each of the five houses is painted a different color, and their inhabitants are of different national extractions, own different pets, drink different beverages and smoke different brands of American cigarettes [sic]. One other thing: in statement 6, right means your right.

— Life International, December 17, 1962

Zebra Puzzle implies some background knowledge, which we must add. First, there is a one-to-one correspondence between houses and colors (and persons, drinks, cigarettes, and pets). Second, we assume that the first house stands on the extreme left and the fifth house stands on the extreme right. The first and the fifth house are not next to each other. A house is on the right of another house if the latter is to the left of the former. Right-to and left-to imply next-to and next-to is a symmetric relation. Next, we list the facts and the rules that encode the puzzle. For clarity statement is related to an appropriate English statement from the puzzle or from the implicit information mentioned above.

Sentences 1-5 provide the *house*-, *color*-, *person*-, *cigarette*-, and *pet*-facts.

1 There are five houses: #1, #2, #3, #4, and #5.
$house(1) : \mathbf{t}.\quad house(2) : \mathbf{t}.\quad house(3) : \mathbf{t}.\quad house(4) : \mathbf{t}.\quad house(5) : \mathbf{t}.$
2 There are five colors: yellow, blue, red, ivory, and green.
$color(yellow) : \mathbf{t}.\quad color(blue) : \mathbf{t}.\quad color(red) : \mathbf{t}.\quad color(ivory) : \mathbf{t}.\quad color(green) : \mathbf{t}.$
3 There are five people: Norwegian, Ukrainian, Englishman, Spaniard, and Japanese.
$person(norwegian) : \mathbf{t}.\quad person(ukrainian) : \mathbf{t}.\quad person(englishman) : \mathbf{t}.$

[7] `https://en.wikipedia.org/wiki/Zebra_Puzzle`
[8] `http://www.cs.miami.edu/~tptp/cgi-bin/SeeTPTP?Category=Problems&Domain=PUZ&File=PUZ010-1.p`

$person(spaniard) : \mathbf{t}.$ $person(japanese) : \mathbf{t}.$

4 There are five drinks: water, tea, milk, orange juice, and coffee.
$drink(water) : \mathbf{t}.$ $drink(tea) : \mathbf{t}.$ $drink(milk) : \mathbf{t}.$ $drink(orange_juice) : \mathbf{t}.$
$drink(coffee) : \mathbf{t}.$

5 There are five cigarettes: Kools, Chesterfield, Old Gold, Luck Strike, and Parliament.
$cigarette(kools) : \mathbf{t}.$ $cigarette(chesterfield) : \mathbf{t}.$ $cigarette(old_gold) : \mathbf{t}.$
$cigarette(lucky_strike) : \mathbf{t}.$ $cigarette(parliament) : \mathbf{t}.$

6 There are five pets: fox, horse, snails, dog, and zebra.
$pet(fox) : \mathbf{t}.$ $pet(horse) : \mathbf{t}.$ $pet(snails) : \mathbf{t}.$ $pet(dog) : \mathbf{t}.$ $pet(zebra) : \mathbf{t}.$

Sentences 7-16 describe the implicit knowledge about the one-to-one correspondence between houses, colors, persons, drinks, cigarettes, and pets. They are encoded as the cardinality constraints on the predicates $house_color(H,C)$, $house_nationality(H,N)$, $house_drink(H,D)$, $house_smoke(H,C)$, and $house_pet(H,P)$. We encode Sentences 7 - 8 using Principles 5 and 6. This will cause the second rule in Sentence 8 to be repeated as part of the encoding of Sentence 7, so we omit the duplicate. The rest of the sentences follow the same idea.

7 Each house has exactly one color.
$1\{house_color(H,C) : \mathbf{t} \ if \ color(C) : \mathbf{t}\}1 \leftarrow house(H) : \mathbf{t}.$
$house_color(H,C) : \mathbf{f} \leftarrow house(H) : \mathbf{t} \wedge color(C) : \mathbf{t} \wedge \neg house_color(H,C) : \mathbf{t}.$

8 Each color is for exactly one house.
$1\{house_color(H,C) : \mathbf{t} \ if \ house(H) : \mathbf{t}\}1 \leftarrow color(C) : \mathbf{t}.$

9 Each house is home for exactly one person.
$1\{house_nationality(H,N) : \mathbf{t} \ if \ person(N) : \mathbf{t}\}1 \leftarrow house(H) : \mathbf{t}.$
$house_nationality(H,N) : \mathbf{f} \leftarrow house(H) : \mathbf{t} \wedge person(N) : \mathbf{t} \wedge \neg house_nationality(H,N) : \mathbf{t}.$

10 Each person lives in exactly one house.
$1\{house_nationality(H,N) : \mathbf{t} \ if \ house(H) : \mathbf{t}\}1 \leftarrow person(N) : \mathbf{t}.$

11 Each house has exactly one favorite drink.
$1\{house_drink(H,D) : \mathbf{t} \ if \ drink(D) : \mathbf{t}\}1 \leftarrow house(H) : \mathbf{t}.$
$house_drink(H,D) : \mathbf{f} \leftarrow house(H) : \mathbf{t} \wedge drink(D) : \mathbf{t} \wedge \neg house_drink(H,D) : \mathbf{t}.$

12 Each drink is drunk in exactly one house.
$1\{house_drink(H,D) : \mathbf{t} \ if \ house(H) : \mathbf{t}\}1 \leftarrow drink(D) : \mathbf{t}.$

13 Each house has exactly one brand of cigarettes.
$1\{house_smoke(H,S) : \mathbf{t} \ if \ cigarette(S) : \mathbf{t}\}1 \leftarrow house(H) : \mathbf{t}.$
$house_smoke(H,S) : \mathbf{f} \leftarrow house(H) : \mathbf{t} \wedge cigarette(S) : \mathbf{t} \wedge \neg house_smoke(H,S) : \mathbf{t}.$

14 Each brand of cigarettes is smoked in exactly one house.
$1\{house_smoke(H,S) : \mathbf{t} \ if \ house(H) : \mathbf{t}\}1 \leftarrow cigarette(S) : \mathbf{t}.$

15 Each house has exactly one pet.
$1\{house_pet(H,P) : \mathbf{t} \ if \ pet(P) : \mathbf{t}\}1 \leftarrow house(H) : \mathbf{t}.$
$house_pet(H,P) : \mathbf{f} \leftarrow house(H) : \mathbf{t} \wedge pet(P) : \mathbf{t} \wedge \neg house_pet(H,P) : \mathbf{t}.$

16 Each pet is kept in exactly one house.
$1\{house_pet(H,P) : \mathbf{t} \ if \ house(H) : \mathbf{t}\}1 \leftarrow pet(P) : \mathbf{t}.$

Sentences 17-20 correspond to Sentences (b)-(e) in the original puzzle. The encoding is based on Principles 1 and 2, where contrapositive inference and propagation of inconsistency are allowed for some literals but not others. Notice that it is undesirable to allow propagation of inconsistency from *house*-facts, since it is unreasonable to conclude that somebody lives in a house

while being unsure that something is a house. Contrapositive reasoning is also inappropriate here. For instance, if there is uncertainty about the color of a house, it is unreasonable to conclude that something is not a house because one is much more likely to discern a house than its color. So, in the next group of rules, we separate $house(H)$ from color and other facts using use ontological implication $\leftarrow$.

17 If the Englishman lives in a house then the color of the house is red.
$(house_color(H, red) : \mathbf{t} \ <\sim \ house_nationality(H, englishman) : \mathbf{t})$
$\leftarrow \ house(H) : \mathbf{t} \ \wedge \ \neg house(H) : \top.$

18 If the Spaniard lives in a house then dog is kept in the house.
$(house_pet(H, dog) : \mathbf{t} \ <\sim \ house_nationality(H, spaniard) : \mathbf{t})$
$\leftarrow \ house(H) : \mathbf{t} \ \wedge \ \neg house(H) : \top.$

19 If the color of a house is green then coffee is drunk in the house.
$(house_drink(H, coffee) : \mathbf{t} \ <\sim \ house_color(H, green) : \mathbf{t})$
$\leftarrow \ house(H) : \mathbf{t} \ \wedge \ \neg house(H) : \top.$

20 If the Ukrainian lives in a house then tea is drunk in the house.
$(house_drink(H, tea) : \mathbf{t} \ <\sim \ house_nationality(H, ukrainian) : \mathbf{t})$
$\leftarrow \ house(H) : \mathbf{t} \ \wedge \ \neg house(H) : \top.$

Sentences 21-22 define the implicit knowledge of the *right(X,Y)* relation. We encode each sentence based on Principle 2.

21 A house numbered X is to the right of another house numbered Y if $\mathrm{X} - 1 = \mathrm{Y}$.
$right(X, Y) : \mathbf{t} \ \leftarrow \ house(X) : \mathbf{t} \ \wedge \ house(Y) : \mathbf{t} \ \wedge \ X - 1 \ = \ Y.$
$\neg house(X) : \top \ \wedge \ \neg house(Y) : \top.$

22 A house numbered X is not to right of another house numbered Y if $\mathrm{X} - 1 \neq \mathrm{Y}$.
$right(X, Y) : \mathbf{f} \ \leftarrow \ house(X) : \mathbf{t} \ \wedge \ house(Y) : \mathbf{t} \ \wedge \ X - 1 \ \neq \ Y.$
$\neg house(X) : \top \ \wedge \ \neg house(Y) : \top.$

Sentences 23-25 correspond to Sentence (f). The encoding of Sentences 23-24 is straightforward. The encoding of Sentence 25 is based on Principles 1 and 2.

23 The color of the first house is not green.
$house_color(1, green) : \mathbf{f}.$

24 The color of the fifth house is not ivory.
$house_color(5, ivory) : \mathbf{f}.$

25 If a house is to the right of another house and the color of the former house is green then the color of the latter house is ivory.
$(house_color(Left, ivory) : \mathbf{t} \ <\sim \ house_color(Right, green) : \mathbf{t})$
$\leftarrow \ house(Left) : \mathbf{t} \ \wedge \ house(Right) : \mathbf{t} \ \wedge \ right(Right, Left) : \mathbf{t} \ \wedge$
$\neg house(Left) : \top \ \wedge \ \neg house(Right) : \top \ \wedge \ \neg right(Right, Left) : \top.$

Sentences 26-27 below correspond to Sentences (g)-(h). The encoding is based on Principles 1 and 2.

26 If Old Gold is smoked in a house then snails are kept in the house.
$(house_pet(H, snails) : \mathbf{t} \ <\sim \ house_smoke(H, old_gold) : \mathbf{t})$
$\leftarrow \ house(H) : \mathbf{t} \ \wedge \ \neg house(H) : \top.$

27 If Kools is smoked in a house then the color of the house is yellow.

$$(house_smoke(H,kools):\mathbf{t} \ <\sim \ house_color(H,yellow):\mathbf{t}) \leftarrow house(H):\mathbf{t} \wedge \neg house(H):\top.$$

Sentences 28-29 below correspond to Sentences (i)-(j). The encoding is straightforward.

28 Milk is drunk in the middle house.
$house_drink(3,milk):\mathbf{t}$.

29 The Norwegian lives in the first house.
$house_nationality(1,norwegian):\mathbf{t}$.

Sentences 30-31 state the implicit knowledge of the *next(X,Y)* relation. We encode each sentence based on Principle 2.

30 A house is next to another house if their house numbers differ by 1
$next(X,Y):\mathbf{t} \leftarrow house(X):\mathbf{t} \wedge house(Y):\mathbf{t} \wedge |X-Y| = 1. \wedge \neg house(X):\top \wedge \neg house(Y):\top.$

31 A house is not next to another house if their house numbers do not differ by 1
$next(X,Y):\mathbf{f} \leftarrow house(X):\mathbf{t} \wedge house(Y):\mathbf{t} \wedge |X-Y| \neq 1. \wedge \neg house(X):\top \wedge \neg house(Y):\top.$

Sentences 32-34 correspond to Sentence (k). Sentences 32-33 are encoded based on Principle 1. Sentence 34 is encoded based on Principles 1 and 2.

32 If Chesterfield is smoked in the first house then fox is kept in the second house.
$house_pet(2,fox):\mathbf{t} \ <\sim \ house_smoke(1,chesterfield):\mathbf{t}$.

33 If Chesterfield is smoked in the fifth house then fox is kept in the fourth house.
$house_pet(4,fox):\mathbf{t} \ <\sim \ house_smoke(5,chesterfield):\mathbf{t}$.

34 If a house is next to another house and Chesterfield is smoked in the former house then fox is kept in the latter house.
$(house_pet(H1,fox):\mathbf{t} \vee house_pet(H3,fox):\mathbf{t} \ <\sim \ house_smoke(H2,chesterfield):\mathbf{t}) \leftarrow house(H1):\mathbf{t} \wedge house(H2):\mathbf{t} \wedge house(H3):\mathbf{t} \wedge H1 \neq H3 \wedge next(H1,H2):\mathbf{t} \wedge next(H2,H3):\mathbf{t} \wedge \neg house(H1):\top \wedge \neg house(H2):\top \wedge \neg house(H3):\mathbf{t}:\top \wedge \neg next(H1,H2):\mathbf{t}:\top \wedge \neg next(H2,H3):\top.$

Sentences 35-37 correspond to Sentence (l). The encoding follows the same idea as Sentences 32-34.

35 If Kools are smoked in the first house then horse is kept in the second house
$house_pet(2,horse):\mathbf{t} \ <\sim \ house_smoke(1,kools):\mathbf{t}$.

36 If Kools are smoked in the fifth house then horse is kept in the fourth house.
$house_pet(4,horse):\mathbf{t} \ <\sim \ house_smoke(5,kools):\mathbf{t}$.

37 If a house is next to another house and Kools are smoked in the former house then horse is kept in the latter house.
$(house_pet(H1,horse):\mathbf{t} \vee house_pet(H3,horse):\mathbf{t} \ <\sim \ house_smoke(H2,kools):\mathbf{t}) \leftarrow house(H1):\mathbf{t} \wedge house(H2):\mathbf{t} \wedge house(H3):\mathbf{t} \wedge H1 \neq H3 \wedge next(H1,H2):\mathbf{t} \wedge next(H2,H3):\mathbf{t} \wedge \neg house(H1):\top \wedge \neg house(H2):\top \wedge \neg house(H3):\mathbf{t}:\top \wedge \neg next(H1,H2):\mathbf{t}:\top \wedge \neg next(H2,H3):\top.$

Sentences 38-39 below correspond to Sentences (m)-(n). The encoding is based on Principles 1 and 2.

38 If Lucky Strike is smoked in a house then orange juice is drunk in that house.
$(house_drink(H, orange_juice) : \mathbf{t} \ <\!\sim\ house_smoke(H, lucky_strike) : \mathbf{t})$
$\leftarrow\ house(H) : \mathbf{t} \wedge \neg house(H) : \top.$

39 If the Japanese lives in a house then parliament is smoked in that house.
$(house_smoke(H, parliament) : \mathbf{t} \ <\!\sim\ house_nationality(H, japanese) : \mathbf{t})$
$\leftarrow\ house(H) : \mathbf{t} \wedge \neg house(H) : \top.$

Sentences 40-42 correspond to Sentence (o). The encoding follows the same idea as Sentences 32-34.

40 If the Norwegian lives in the first house then the color of the second house is blue.
$house_color(2, blue) : \mathbf{t} \ <\!\sim\ house_nationality(1, norwegian) : \mathbf{t}.$

41 If the Norwegian lives in the fifth house then the color of the fourth house is blue.
$house_color(4, blue) : \mathbf{t} \ <\!\sim\ house_nationality(5, norwegian) : \mathbf{t}.$

42 If a house is next to another house and the Norwegian lives in the former house then the color of the latter house is blue.
$(house_color(H1, blue) : \mathbf{t} \vee house_color(H3, blue) : \mathbf{t}$
$<\!\sim\ house_nationality(H2, norwegian) : \mathbf{t})$
$\leftarrow\ house(H1) : \mathbf{t} \wedge house(H2) : \mathbf{t} \wedge house(H3) : \mathbf{t} \wedge$
$H1 \neq H3 \wedge next(H1, H2) : \mathbf{t} \wedge next(H2, H3) : \mathbf{t} \wedge$
$\neg house(H1) : \top \wedge \neg house(H2) : \top \wedge \neg house(H3) : \mathbf{t} : \top \wedge$
$\neg next(H1, H2) : \mathbf{t} : \top \wedge \neg next(H2, H3) : \top.$

Next we define the consistency preference relation $<_{\mathscr{S}}$, where $\mathscr{S} = (s_1, s_2, \mathscr{B}_\top)$, which implements Principle 4. Here s_1 says that we hold greater confidence in the information about the *house-*, *color-*, *person-*, *cigarette-*, and *pet-*, *next*, *right*-facts. We remind that "greater confidence" here means that these facts are least likely to be inconsistent. The preference s_2 says that next we are likely to be confident in the information given by the $house_color(H, C)$, $house_nationality(H, N)$, $house_drink(H, D)$, $house_smoke(H, C)$, $house_pet(H, P)$ facts. The last component in $\mathscr{S}$, $\mathscr{B}_\top$, is the usual default that gives preference to the most e-consistent models.

$$\begin{aligned}
s_1 = \{&house(1) : \top,\ \ldots,\ house(5) : \top,\ color(yellow) : \top,\ \ldots,\ color(green) : \top,\\
&person(norwegian) : \top,\ \ldots,\ person(japanese) : \top,\ drink(water) : \top,\ \ldots,\\
&drink(coffe) : \top,\ cigarette(kools) : \top,\ \ldots,\ cigarette(parliament) : \top,\\
&pet(fox) : \top,\ \ldots,\ pet(zebra) : \top\}.\\
s_2 = \{&house_color(1, yellow) : \top,\ \ldots,\ house_color(5, green) : \top,\\
&house_nationality(1, norwegian) : \top,\ \ldots,\ house_nationality(5, japanese) : \top,\\
&house_drink(1, water) : \top,\ \ldots,\ house_drink(5, coffe) : \top,\\
&house_smoke(1, kools) : \top,\ \ldots,\ house_smoke(5, parliament) : \top,\\
&house_pet(1, fox) : \top,\ \ldots,\ house_pet(5, zebra) : \top\}.
\end{aligned}$$

To better illustrate the result, we define a single predicate tuple/6 to represent the combined information of a house: its associated color, person, drink, cigarette, and pet.

$$\begin{aligned}
tuple(H, C, N, D, S, P) : \mathbf{t} \leftarrow\ &house_color(H, C) : \mathbf{t} \wedge house_nationality(H, N) : \mathbf{t} \wedge\\
&house_drink(H, D) : \mathbf{t} \wedge house_smoke(H, S) : \mathbf{t} \wedge\\
&house_pet(H, P) : \mathbf{t}.
\end{aligned}$$

There is a single most consistency preferred model:

$tuple(1, yellow, norwegian, water, kools, fox) : \mathbf{t}$
$tuple(2, blue, ukrainian, tea, chesterfield, horse) : \mathbf{t}$
$tuple(3, red, englishman, milk, old_gold, snails) : \mathbf{t}$
$tuple(4, ivory, spaniard, orange_juice, lucky_strike, dog) : \mathbf{t}$
$tuple(5, green, japanese, coffee, parliament, zebra) : \mathbf{t}$

Not surprisingly, this solution matches the usual correct solution because so far we have not injected any inconsistency and so, for example, the preferences s_1 and s_2 play no role.

Next, we illustrate three cases of injection of inconsistency into the puzzle. Since complete models tend to be rather large, we show only $tuple/6$ and $\top$-predicates.

43 The Ukrainian lives in the middle house — Variation 1
$house_nationality(3, ukrainian) : \mathbf{t}$.

There are two models:

$m_1 = \{tuple(1, yellow, norwegian, water, kools, zebra) : \mathbf{t}$
$tuple(2, blue, japanese, tea, parliament, horse) : \mathbf{t}$
$tuple(3, ivory, ukrainian, milk, old_gold, snails) : \mathbf{t}$
$tuple(4, green, spaniard, coffee, chesterfield, dog) : \mathbf{t}$
$tuple(5, red, englishman, orange_juice, lucky_strike, fox) : \mathbf{t}$
$house_nationality(3, ukrainian) : \top\}$

$m_2 = \{tuple(1, yellow, norwegian, tea, kools, zebra) : \mathbf{t}$
$tuple(2, blue, japanese, water, parliament, horse) : \mathbf{t}$
$tuple(3, ivory, ukrainian, milk, old_gold, snails) : \mathbf{t}$
$tuple(4, green, spaniard, coffee, chesterfield, dog) : \mathbf{t}$
$tuple(5, red, englishman, orange_juice, lucky_strike, fox) : \mathbf{t}$
$house_nationality(3, ukrainian) : \top\}$

The puzzle originally implied that the Ukrainian lives in the second house. Therefore, Variation 1 generates an inconsistency about the Ukrainian being in the middle house. By Sentence 20, we cannot derive $house_drink(3, tea) : \mathbf{t}$ because propagation of inconsistency is blocked in this case. This sanctions two possibilities: one where $house_drink(1, water) : \mathbf{t}$ and $house_drink(2, tea) : \mathbf{t}$ hold and the other where $house_drink(1, tea) : \mathbf{t}$ and $house_drink(2, water) : \mathbf{t}$ hold.

44 The Lucky Strike is smoked in the middle house — Variation 2
$house_smoke(3, lucky_strike) : \mathbf{t}$.

Again, we have two models:

$m_1 = \{tuple(1, yellow, norwegian, water, kools, fox) : \mathbf{t}$
$tuple(2, blue, ukrainian, tea, chesterfield, horse) : \mathbf{t}$
$tuple(3, ivory, spaniard, milk, lucky_strike, dog) : \mathbf{t}$
$tuple(4, green, japanese, coffee, parliament, zebra) : \mathbf{t}$
$tuple(5, red, englishman, orange_juice, old_gold, snails) : \mathbf{t}$
$house_smoke(3, lucky_strike) : \top\}$

$m_2 = \{tuple(1, yellow, norwegian, orange_juice, kools, fox) : \mathbf{t}$

$tuple(2,\ blue,\ ukrainian,\ tea,\ chesterfield,\ horse)$: **t**
$tuple(3,\ ivory,\ spaniard,\ milk,\ lucky_strike,\ dog)$: **t**
$tuple(4,\ green,\ japanese,\ coffee,\ parliament,\ zebra)$: **t**
$tuple(5,\ red,\ englishman,\ water,\ old_gold,\ snails)$: **t**
$house_smoke(3,\ lucky_strike)$: $\top\}$

The puzzle originally implied that the Lucky Strike is smoked in the fourth house. As a result, Variation 2 generates an inconsistency regarding the Lucky Strike being smoked in the middle house. By Sentence 38, we cannot derive $house_drink(3, orange_juice)$: **t** because propagation of inconsistency is blocked by the epistemic implication. This sanctions two possibilities: one where $house_drink(1, water)$: **t** and $house_drink(5, orange_juice)$: **t** hold and the other where $house_drink(1, orange_juice)$: **t** and $house_drink(5, water)$: **t** hold.

45 Milk is not drunk in the middle house — Variation 3
$house_drink(3, milk)$: **f**.

Variation 3 generates the same model as Zebra Puzzle's original solution except for an additional inconsistent fact $house_drink(3, milk)$: $\top$. This is because the puzzle originally implied that milk *is* drunk in the middle house.

Appendix C Marathon Puzzle in APC_{LP} with Inconsistency Injections

We present here a complete APC_{LP} encoding of Marathon Puzzle (C. Guéret and Sevaux 2000; Schwitter 2012). As with the previous puzzles, the encoding highlights the principles of Section 6 and we also discuss several cases of inconsistency injection.

Marathon puzzle is as follows:

> Dominique, Ignace, Naren, Olivier, Philippe, and Pascal have arrived as the first six at the Paris marathon.
> Olivier has not arrived last.
> Dominique, Pascal and Ignace have arrived before Naren and Olivier.
> Dominique who was third last year has improved this year.
> Philippe is among the first four.
> Ignace has arrived neither in second nor third position.
> Pascal has beaten Naren by three positions.
> Neither Ignace nor Dominique are in the fourth position.

The original description implies some implicit background knowledge. First, no runners arrive at the same time (i.e., each runner has a unique arrival position and vice versa). Second, a runner arrives before another runner if the first runner's position number is smaller than the second runner's position.

There is a unique solution for Marathon Puzzle where Ignace arrives first, followed by Dominique, Pascal, Philippe, Olivier, and Naren, in that order. Next, we show the encoding of Marathon puzzle.

1 Dominique, Ignace, Naren, Olivier, Philippe, and Pascal have arrived as the first six at the Paris marathon.
$runner(dominique)$: **t**. $runner(naren)$: **t**. $runner(ignace)$: **t**. $runner(olivier)$: **t**. $runner(philippe)$: **t**. $runner(pascal)$: **t**.

$position(1):\mathbf{t}.\quad position(2):\mathbf{t}.\quad position(3):\mathbf{t}.\quad position(4):\mathbf{t}.\quad position(5):\mathbf{t}.$
$position(6):\mathbf{t}.$

Sentences 2 and 3 are encoded based on Principles 5 and 6. This will cause the second rule in Sentence 3 to be repeated as part of the encoding of Sentence 2 so, as before, we omit the duplicate.

2 Every runner has exactly one position.
$1\{has_position(R,P):\mathbf{t} \quad if \quad position(P):\mathbf{t}\}1 \leftarrow runner(R):\mathbf{t}.$
$has_position(R,P):\mathbf{f} \leftarrow runner(R):\mathbf{t} \wedge position(P):\mathbf{t} \wedge \neg has_position(R,P):\mathbf{t}.$
3 Every position belongs to exactly one runner.
$1\{has_position(R,P):\mathbf{t} \quad if \quad runner(R):\mathbf{t}\}1 \leftarrow position(P):\mathbf{t}.$

Encoding of Sentence 4 is straightforward.

4 Olivier has not arrived last.
$has_position(olivier,6):\mathbf{f}.$

Sentence 5 is encoded by three rules. The first and second rules are the result of applying Principle 2 defining the *before*-relation based on the position of runners. The third rule represents complete knowledge of the *before*-relation based on Principle 5.

5 If a runner R1 has a position P1 and another runner R2 has a position P2 and P1 $<$ P2 then R1 is *before* R2.

$$\begin{aligned}
before(R1,R2):\mathbf{t} \leftarrow\ & runner(R1):\mathbf{t} \wedge position(P1):\mathbf{t} \wedge has_position(R1,P1):\mathbf{t} \wedge \\
& runner(R2):\mathbf{t} \wedge position(P2):\mathbf{t} \wedge has_position(R2,P2):\mathbf{t} \wedge \\
& P1 < P2 \wedge \neg runner(R1):\top \wedge \neg runner(R2):\top \wedge \\
& \neg position(P1):\top \wedge \neg position(P2):\top \wedge \\
& \neg has_position(R1,P1):\top \wedge \neg has_position(R2,P2):\top.
\end{aligned}$$

$$\begin{aligned}
before(R2,R1):\mathbf{f} \leftarrow\ & runner(R1):\mathbf{t} \wedge position(P1):\mathbf{t} \wedge has_position(R1,P1):\mathbf{t} \wedge \\
& runner(R2):\mathbf{t} \wedge position(P2):\mathbf{t} \wedge has_position(R2,P2):\mathbf{t} \wedge \\
& P1 < P2 \wedge \neg runner(R1):\top \wedge \neg runner(R2):\top \wedge \\
& \neg position(P1):\top \wedge \neg position(P2):\top \wedge \\
& \neg has_position(R1,P1):\top \wedge \neg has_position(R2,P2):\top.
\end{aligned}$$

$before(R1,R2):\mathbf{f} \leftarrow runner(R1):\mathbf{t} \wedge runner(R2):\mathbf{t} \wedge \neg before(R1,R2):\mathbf{t}.$

Encoding of Sentences 6 - 9 is straightforward.

6 Dominique, Pascal and Ignace have arrived before Naren and Olivier.
$before(dominique,naren):\mathbf{t}.\quad before(pascal,naren):\mathbf{t}.\quad before(ignace,naren):\mathbf{t}.$
$before(dominique,olivier):\mathbf{t}.\quad before(pascal,olivier):\mathbf{t}.\quad before(ignace,olivier):\mathbf{t}.$
7 Dominique who was third last year has improved this year.
$has_position(dominique,1):\mathbf{t} \vee has_position(dominique,2):\mathbf{t}.$
8 Philippe is among the first four.
$has_position(philippe,1):\mathbf{t} \vee has_position(philippe,2):\mathbf{t} \vee$
$has_position(philippe,3):\mathbf{t} \vee has_position(philippe,4):\mathbf{t}.$
9 Ignace has arrived neither in second nor third position.
$has_position(ignace,2):\mathbf{f}.$
$has_position(ignace,3):\mathbf{f}.$

Encoding of Sentence 10 is based on Principle 2. The first rule encodes that if the position of Pascal is known, then the position of Naren is the position of Pascal plus 3. The second rule encodes that if the position of Naren is known, then the position of Pascal is the position of Naren minus 3.

10 Pascal has beaten Naren by three positions.
$$\begin{aligned}&\mathit{has_position}(\mathit{naren},P2):\mathbf{t} \leftarrow \mathit{has_position}(\mathit{pascal},P1):\mathbf{t} \wedge \mathit{position}(P1):\mathbf{t} \wedge \\ &\qquad \mathit{position}(P2):\mathbf{t} \wedge P2 = P1+3 \wedge \neg\mathit{position}(P1):\top \wedge \\ &\qquad \neg\mathit{position}(P2):\top \wedge \neg\mathit{has_position}(\mathit{pascal},P1):\top.\\ &\mathit{has_position}(\mathit{pascal},P1):\mathbf{t} \leftarrow \mathit{has_position}(\mathit{naren},P2):\mathbf{t} \wedge \mathit{position}(P1):\mathbf{t} \wedge \\ &\qquad \mathit{position}(P2):\mathbf{t} \wedge P2 = P1+3 \wedge \neg\mathit{position}(P1):\top \wedge \\ &\qquad \neg\mathit{position}(P2):\top \wedge \neg\mathit{has_position}(\mathit{naren},P2):\top.\end{aligned}$$

Encoding of Sentence 11 is straightforward.

11 Neither Ignace nor Dominique are in the fourth position.
$\mathit{has_position}(\mathit{ignace},4):\mathbf{f}.\quad \mathit{has_position}(\mathit{dominique},4):\mathbf{f}.$

Next we define the consistency preference relation $<_{\mathscr{S}}$, where $\mathscr{S} = (s_1, s_2, \mathscr{B}_\top)$, which implements Principle 4. Here s_1 says that we hold greater confidence in the information about the *runner*- and *position*-facts and therefore these types of facts are less likely to be inconsistent. The preference s_2 says that next we are likely to be confident in the information given by the *has_position*- and *before*-facts. The last component in $\mathscr{S}$, $\mathscr{B}_\top$, is the usual default that gives preference to the most e-consistent models.

$$\begin{aligned}s_1 = \{&\mathit{runner}(\mathit{dominique}):\top,\ \mathit{runner}(\mathit{naren}):\top,\ \mathit{runner}(\mathit{ignace}):\top,\\ &\mathit{runner}(\mathit{olivier}):\top,\ \mathit{runner}(\mathit{philippe}):\top,\ \mathit{runner}(\mathit{pascal}):\top,\\ &\mathit{position}(1):\top,\ \ldots,\ \mathit{position}(6):\top\}.\\ s_2 = \{&\mathit{has_position}(\mathit{dominique},1):\top,\ \ldots,\ \mathit{has_position}(\mathit{dominique},6):\top,\ \ldots\\ &\mathit{has_position}(\mathit{pascal},1):\top,\ \ldots,\ \mathit{has_position}(\mathit{pascal},6):\top,\\ &\mathit{before}(\mathit{dominique},\mathit{naren}):\top,\ \ldots,\ \mathit{before}(\mathit{dominique},\mathit{pascal}):\top,\ \ldots\\ &\mathit{before}(\mathit{pascal},\mathit{dominique}):\top,\ \ldots,\ \mathit{before}(\mathit{pascal},\mathit{philippe}):\top\}.\end{aligned}$$

There is a single most consistency-preferred model as expected. The puzzle, as stated, has no inconsistent information so we show only the *has_position*-facts.

$$\begin{aligned}m = \{&\mathit{has_position}(\mathit{ignace},1):\mathbf{t},\ \mathit{has_position}(\mathit{dominique},2):\mathbf{t},\\ &\mathit{has_position}(\mathit{pascal},3):\mathbf{t},\ \mathit{has_position}(\mathit{philippe},4):\mathbf{t},\\ &\mathit{has_position}(\mathit{olivier},5):\mathbf{t},\ \mathit{has_position}(\mathit{naren},6):\mathbf{t}\}.\end{aligned}$$

In addition, there are many *before*-facts which encode the sequential order of arrival of the runners.

Next, we illustrate several cases of injection of inconsistency into the puzzle. Since complete models tend to be rather large, we show only the *has_position*- and $\top$-predicates.

12 Pascal arrives in the sixth position — Variation 1
$\mathit{has_position}(\mathit{pascal},6):\mathbf{t}.$

There is one model and it contains inconsistencies:

$$\begin{aligned}m = \{&\mathit{has_position}(\mathit{ignace},1):\mathbf{t},\ \mathit{has_position}(\mathit{dominique},2):\mathbf{t},\\ &\mathit{has_position}(\mathit{naren},3):\mathbf{t},\ \mathit{has_position}(\mathit{philippe},4):\mathbf{t},\\ &\mathit{has_position}(\mathit{olivier},5):\mathbf{t},\ \mathit{has_position}(\mathit{pascal},6):\mathbf{t},\end{aligned}$$

$\mathit{before}(\mathit{pascal},\mathit{naren}) : \top,\ \mathit{before}(\mathit{pascal},\mathit{olivier}) : \top\}$.

The *before*-facts still represent the sequential order of arrival, so $\mathit{before}(\mathit{pascal},\mathit{naren}) : \mathbf{t}$ and $\mathit{before}(\mathit{pascal},\mathit{olivier}) : \mathbf{t}$ are true in the model. However, Sentence 12 contradicts these facts, so $\mathit{before}(\mathit{pascal},\mathit{naren}) : \mathbf{f}$ and $\mathit{before}(\mathit{pascal},\mathit{olivier}) : \mathbf{f}$ are also true. Therefore, these facts become inconsistent. Sentence 10 does not imply any inconsistencies beyond the ones already mentioned.

13 Ignace arrives in the second position. — Variation 2
$\mathit{has_position}(\mathit{ignace},2) : \mathbf{t}$.

There is one model:
$m = \{\mathit{has_position}(\mathit{dominique},1) : \mathbf{t},\ \mathit{has_position}(\mathit{ignace},2) : \top,$
$\mathit{has_position}(\mathit{pascal},3) : \mathbf{t},\ \mathit{has_position}(\mathit{philippe},4) : \mathbf{t},$
$\mathit{has_position}(\mathit{olivier},5) : \mathbf{t},\ \mathit{has_position}(\mathit{naren},6) : \mathbf{t}\}$.

Given that Sentence 13 contradicts Sentence 9, the fact $\mathit{has_position}(\mathit{ignace},2)$ becomes inconsistent and the order of arrival of Dominique and Ignace are swapped.

14 Philippe arrives before Dominique. — Variation 3
$\mathit{before}(\mathit{philippe},\mathit{dominique}) : \mathbf{t}$.

Now, we get 4 models:
$m_1 = \{\mathit{has_position}(\mathit{ignace},1) : \mathbf{t},\ \mathit{has_position}(\mathit{dominique},2) : \mathbf{t},$
$\mathit{has_position}(\mathit{pascal},3) : \mathbf{t},\ \mathit{has_position}(\mathit{philippe},4) : \mathbf{t},$
$\mathit{has_position}(\mathit{olivier},5) : \mathbf{t},\ \mathit{has_position}(\mathit{naren},6) : \mathbf{t},$
$\mathit{before}(\mathit{philippe},\mathit{dominique}) : \top\}$.
$m_2 = \{\mathit{has_position}(\mathit{philippe},1) : \mathbf{t},\ \mathit{has_position}(\mathit{dominique},2) : \mathbf{t},$
$\mathit{has_position}(\mathit{pascal},3) : \mathbf{t},\ \mathit{has_position}(\mathit{ignace},4) : \top,$
$\mathit{has_position}(\mathit{olivier},5) : \mathbf{t},\ \mathit{has_position}(\mathit{naren},6) : \mathbf{t}\}$.
$m_3 = \{\mathit{has_position}(\mathit{philippe},1) : \mathbf{t},\ \mathit{has_position}(\mathit{dominique},2) : \mathbf{t},$
$\mathit{has_position}(\mathit{pascal},3) : \mathbf{t},\ \mathit{has_position}(\mathit{olivier},4) : \mathbf{t},$
$\mathit{has_position}(\mathit{ignace},5) : \mathbf{t},\ \mathit{has_position}(\mathit{naren},6) : \mathbf{t},$
$\mathit{before}(\mathit{ignace},\mathit{olivier}) : \top\}$.
$m_4 = \{\mathit{has_position}(\mathit{philippe},1) : \mathbf{t},\ \mathit{has_position}(\mathit{dominique},2) : \mathbf{t},$
$\mathit{has_position}(\mathit{naren},3) : \mathbf{t},\ \mathit{has_position}(\mathit{pascal},4) : \mathbf{t},$
$\mathit{has_position}(\mathit{ignace},5) : \mathbf{t},\ \mathit{has_position}(\mathit{olivier},6) : \top,$
$\mathit{before}(\mathit{pascal},\mathit{naren}) : \top,\ \mathit{before}(\mathit{ignace},\mathit{naren}) : \top\}$.

The first model generates the same solution as the original puzzle except for the inconsistency where Philippe arrives before Dominique. The second model places Ignace in the fourth position, which contradicts Sentence 11, so placing Ignace in the fourth position becomes inconsistent. The third model places Olivier before Ignace, which is in contradiction with Sentence 6. All the other constraints are satisfied, so no more inconsistencies are derived. The fourth model places Naren before Pascal and Ignace, which contradicts Sentence 6. Besides, it places Olivier the last, which contradicts Sentence 4.